

CORRECTED VERSION

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
28 March 2002 (28.03.2002)

PCT

(10) International Publication Number  
**WO 02/025588 A3**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**,  
G06T 7/00

(21) International Application Number: PCT/US01/29573

(22) International Filing Date:  
21 September 2001 (21.09.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/234,108 21 September 2000 (21.09.2000) US  
60/234,435 21 September 2000 (21.09.2000) US  
60/234,114 21 September 2000 (21.09.2000) US  
60/234,115 21 September 2000 (21.09.2000) US

(71) Applicant (for all designated States except US): **MD ON-LINE INC.** [US/US]; 99 Hayden Avenue, Lexington, MA 02421 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **STEFANESCU,**

**Dan** [US/US]; 170 Thornton Road, Needham, MA 02492 (US). **LEVENTON, Michael** [US/US]; 60 Baker Avenue, Lexington, MA 02421 (US).


(74) Agents: **MAZZARESE, Robert, A.** et al.; Ropes & Gray, Patent Group, One International Place, Boston, MA 02110 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: MEDICAL IMAGE PROCESSING SYSTEMS



(57) Abstract: The systems described herein employ a pre-processing strategy for image databases to obtain derivative information from images, in particular from within regions of interest (if any) in each image. A similar pre-processing strategy may be applied to a query image so that rapid and highly accurate comparisons may be performed between the query image and images from the image database, based upon the derivative information. The pre-processing strategy may include, for example, normalization of image data, sampling of normalized image data, segmentation of regions of interest, and extraction of feature vectors from the normalized, sampled, segmented image data. The overall strategy may provide for efficacious comparison of new and archived images. Further disclosed herein are techniques employing these strategies in modeling normalized image data for various applications, in organizing and searching databases, in providing a user interface for applying the systems in a medical workplace, and for deploying applications based upon these strategies locally and over computer networks.



WO 02/025588 A3



**Published:**

— *with international search report*

**(88) Date of publication of the international search report:**

28 August 2003

**(48) Date of publication of this corrected version:**

6 November 2003

**(15) Information about Correction:**

see PCT Gazette No. 45/2003 of 6 November 2003, Section II

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## MEDICAL IMAGE PROCESSING SYSTEMS

### Technical Field

The invention relates generally to medical image processing, and more particularly to the acquisition, searching, organization, comparison, and evaluation of medical images.

### Background Art

Significant progress has been made in the capture of medical images, such as CT scans, x-ray images, ultrasound images, magnetic resonance imaging (MRI) and so forth. At the same time, computer-assisted evaluation of images has moved forward with a number of different computational techniques, further aided by continuing improvements in computer processing power and data networking.

Nonetheless, practical applications for computer-assisted evaluation of medical images have proved elusive. For example, in neural radiology, a single three-dimensional, neural MR study may contain large amounts of data, and a database of such studies may consume terabytes of storage or more. Furthermore, from an image-processing point of view, this image-based data is characterized by significant complexity that makes accurate matching a difficult and computationally expensive task, for traditional, state-of-the-art technology. Thus, while processing and matching strategies may be devised for such images using known techniques, these techniques have thus far failed to produce systems suitable for deployment on typical desktop computers and network connections.

There remains a need for systems that provide management and evaluation tools for large-scale image databases, such as medical image databases.

### Disclosure of Invention

The systems described herein employ a pre-processing strategy for image databases to obtain derivative information from images, in particular from within regions of interest (if any) in each image. A similar pre-processing strategy may be applied to a query image so that rapid and highly accurate comparisons may be performed between the query image and images from the image database, based

upon the derivative information. The pre-processing strategy may include, for example, normalization of image data, sampling of normalized image data, segmentation of regions of interest, and extraction of feature vectors from the normalized, sampled, segmented image data. The overall strategy may provide for  
5 efficacious comparison of new and archived images. Further disclosed herein are techniques for modeling normalized image data for various applications, architectures for deploying the systems described herein locally and over computer networks, and a user interface for a medical image processing system.

The invention disclosed herein relates to methods and systems for database  
10 organization and workflow management for medical imaging applications.

According to one aspect of the invention, a method for medical image processing and database organization includes receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique and registering each one of the plurality of  
15 images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of registered images.

In one embodiment, the method further includes receiving a query image, the query image including an instance of a human body part obtained through the  
20 medical imaging technique, registering the query image to the coordinate system in a non-rigid manner to obtain a registered query image, and identifying one or more regions of interest in the registered query image. One or more feature vectors are extracted from the one or more regions of interest in the registered query image, and the query image is matched to one or more of the plurality of images by comparing  
25 the one or more feature vectors from the one or more regions of interest in the registered query image to one or more feature vectors of one or more of the plurality of registered images.

In another embodiment, each one of the plurality of registered images are labeled with a label according to an observed characteristic, and one or more  
30 feature vectors are obtained from each one of the plurality of registered images,

wherein a model is trained to associate the one or more feature vectors with the label.

In yet another embodiment, the method further includes receiving a header for each one of the plurality of images that includes data associated with the one of the plurality of images; obtaining one or more feature vectors from each one of the plurality of registered images; training a model to associate the one or more feature vectors with the header; and applying the model to identify the presence of any errors in a new header for a new image. Alternatively or in addition, the method may also include associating a pathology with each one of the plurality of images; training a model to associate the pathology associated with each image with the one or more feature vectors for that image; and applying the model to identify the presence of the pathology in a new image. Moreover, a spatial probability map of locations of the pathology can be generated from the plurality of registered images and the regions of interest.

The invention also provides for systems and computer programs with computer executable code for performing the methods of the invention.

Embodiments of the invention may also include one or more of the following features. The feature set may be normalized so that each feature and each pathology has predetermined relationship to the output. One or more regions of interest may be identified in each the images, and one or more feature vectors may be extracted from inside the registered query image and from outside regions of interest in the registered query image. The regions of interest in the registered query image may represent the entire registered query image. The query image can be matched to at least some of the images by comparing one or more feature vectors from inside the regions of interest and one of the feature vectors from outside the regions of interest to one or more feature vectors of the registered images. The matching process may be done automatically. A score may be assigned to indicate a similarity between the query image and each one of the plurality of images.

The systems and methods may further include identifying a relative z-position in each one of the plurality of registered images, the relative z-position

corresponding to a position of a planar image along a z-axis, wherein the position can be determined by training a model. Moreover, the plurality of registered images may be filtered according to one or more filtering criteria to obtain a subset of the plurality of registered images to match to the registered query image. The query image may be associated with a pathology based upon pathology data of registered images and/or a match with feature vectors. The images can include magnetic resonance images and/or computerized tomography images obtained from the head, neck, spine, muscles and/or skeleton of a patient. Registration and matching of images may include normalizing, masking, scaling, rotating, and/or warping. The registered images may also be divided into a plurality of discrete tiles and sampled within the discrete tiles.

Other advantageous embodiments may include generating a label for additional images, wherein each label represents a region of interest and/or a position of an image on a z-axis perpendicular to an image plane of the image. Alternatively or in addition, a label may indicate the age, sex, a diagnosis, a presence of contrast agents, an image type, or a diagnostic significance of a region of interest.

The model may be trained by a statistical learning methodology, such as a linear regression model, with one or more coefficients of the model being determined by a least squares fit. The various databases may be organized by locating images, registering images, and labeling images within each one of the databases. Different databases can be labeled with different models, whereby the various databases can be searched using the labels to evaluate a similarity of a query to one or more records in each of the databases.

According to another aspect of the invention, in a radiology workflow management system and method, medical image studies are displayed for review in a patient data workspace. A user-navigable, three-dimensional map of one or more aspects of human physiology relating to the medical image studies are displayed in an atlas workspace. A reference workspace can include textual reference materials relating to the image studies, and a searchable database of one or more studies

showing pathologies described in the textual reference materials. A results workspace displays images matched to at least one of the medical image studies for review, a user-identified pathology, data from the searchable database, or a location specified in the atlas. In response to user input, at least one of the patient data workspace, the atlas workspace, the reference workspace, and the results workspace is selected as a current workspace for display and for receipt of user inputs.

A number of search modes can be provided, wherein a database of images can be searched for one or more images similar to a query image; and/or the database of images can be searched for one or more images matching a location in an atlas which renders a three-dimensional representation of an object; and/or the database of images is searched for a diagnosis that has associated therewith images similar to the query image, the diagnosis and a similarity score being provided for each image retrieved from the database of images. At least some of the search modes can employ a matching system in which a database of medical image studies is pre-processed to non-rigidly register one or more images of the medical image studies in the database to a global coordinate system to provide a plurality of registered medical images so as to extract features from the registered medical images for comparing the registered medical images with the registered query images.

According to another aspect of the invention, a method of assembling a group of at least two entities having a common interest includes comparing data provided by a first entity to stored reference data, establishing a first correspondence between the data provided by the first entity and the reference data; comparing data provided by at least one other entity to the stored reference data; and establishing at least one other correspondence between the data provided by the at least one other entity and the reference data. The method further includes evaluating a similarity measure between the first correspondence and the at least one other correspondence, and forming the group of the first entity and the at least one other entity if the similarity measure exceeds a predetermined threshold.

#### **Brief Description Of Drawings**

The foregoing and other objects and advantages of the invention will be appreciated more fully from the following further description thereof, with reference to the accompanying drawings, wherein:

Fig. 1 shows a schematic diagram of the entities involved in an embodiment  
5 of a method and system disclosed herein;

Fig. 2 shows a block diagram of a server that may be used with the systems described herein;

Fig. 3 shows a page that may be used as a user interface;

Fig. 4 shows a patient workspace of a user interface;

10 Fig. 5 shows an atlas workspace of a user interface;

Fig. 6 shows a reference workspace of a user interface;

Fig. 7 shows a results workspace of a user interface;

Fig. 8 is a flow chart showing a process for processing images according to the systems described herein;

15 Fig. 9 shows several possible arrangements of tiles for sampling a magnetic resonance image;

Fig. 10 is a flowchart of a process for organizing databases;

Fig. 11 shows a state diagram for a workflow management system;

Fig. 12 shows schematically a parameter space of a database and of users  
20 having overlapping interests;

Fig. 13 shows a process flow for entering user data to form/join a group;

Fig. 14 depicts exemplary data entries in user data sheet;

Fig. 15 depicts handling of user data without a match in the database; and

Fig. 16 shows an exemplary process for matching patient data with clinical  
25 trials.

### **Best Mode(s) for Carrying Out the Invention**

To provide an overall understanding of the invention, certain illustrative embodiments will now be described, including a client/server architecture for a medical image processing system. However, it will be understood that the methods  
30 and systems described herein can be suitably adapted to any environment where



image data or other complex data structures are to be organized into a database for modeling, searching, or other further processing, and may be deployed, for example, as a stand-alone desktop computer application, within a corporate intranet or over a private network. Applications of the organized database may include, for example, medical diagnosis, statistical analysis, or preparation of instructional, academic software. These and other applications of the systems described herein are intended to fall within the scope of the invention. More generally, the principles of the invention are generally applicable to any environment where organization or analysis of image data or other complex data structures is desired.

10        Figure 1 shows a schematic diagram of the entities involved in an embodiment of a method and system disclosed herein. In a system 100, a plurality of clients 102, servers 104, and providers 108 are connected via an internetwork 110. It should be understood that any number of clients 102, servers 104, and providers 108 could participate in such a system 100. The system may further include one or more local area networks ("LAN") 112 interconnecting clients 102 through a hub 114 (in, for example, a peer network such as Ethernet) or a local area network server 114 (in, for example, a client-server network). The LAN 112 may be connected to the internetwork 110 through a gateway 116, which provides security to the LAN 112 and ensures operating compatibility between the LAN 112 and the internetwork 110. Any data network may be used as the internetwork 110 and the LAN 112.

25        In one embodiment, the internetwork 110 is the Internet, and the World Wide Web provides a system for interconnecting clients 102 and servers 104 through the Internet 110. The internetwork 110 may include a cable network, a wireless network, and any other networks for interconnecting clients, servers and other devices.

30        An exemplary client 102 includes the conventional components of a client system, such as a processor, a memory (e.g. RAM), a bus which couples the processor and the memory, a mass storage device (e.g. a magnetic hard disk or an optical storage disk) coupled to the processor and the memory through an I/O

controller, and a network interface coupled to the processor and the memory, such as modem, digital subscriber line ("DSL") card, cable modem, network interface card, wireless network card, or other interface device capable of wired, fiber optic, or wireless data communications. One example of such a client 102 is a personal computer equipped with an operating system such as Microsoft Windows 2000, Microsoft Windows NT, Unix, Linux, and Linux variants, along with software support for Internet communication protocols. The personal computer may also include a browser program, such as Microsoft Internet Explorer or Netscape Navigator, to provide a user interface for access to the Internet 110. Although the personal computer is a typical client 102, the client 102 may also be a workstation, mobile computer, Web phone, television set-top box, interactive kiosk, personal digital assistant, or other device capable of communicating over the Internet 110. As used herein, the term "client" is intended to refer to any of the above-described clients 102, as well as proprietary network clients designed specifically for the medical image processing systems described herein, and the term "browser" is intended to refer to any of the above browser programs or other software or firmware providing a user interface for navigating the Internet 110 and/or communicating with the medical image processing systems.

An exemplary server 104 includes a processor, a memory (e.g. RAM), a bus which couples the processor and the memory, a mass storage device (e.g. a magnetic or optical disk) coupled to the processor and the memory through an I/O controller, and a network interface coupled to the processor and the memory. Servers may be organized as layers of clusters in order to handle more client traffic, and may include separate servers for different functions such as a database server, a file server, an application server, and a Web presentation server. Such servers may further include one or more mass storage devices such as a disk farm or a redundant array of independent disk ("RAID") system for additional storage and data integrity. Read-only devices, such as compact disc drives and digital versatile disc drives, may also be connected to the servers. Suitable servers and mass storage devices are manufactured by, for example, Compaq, IBM, and Sun Microsystems. As used

herein, the term "server" is intended to refer to any of the above-described servers 104.

Focusing now on the internetwork 110, one embodiment is the Internet. The structure of the Internet 110 is well known to those of ordinary skill in the art and includes a network backbone with networks branching from the backbone. These branches, in turn, have networks branching from them, and so on. The backbone and branches are connected by routers, bridges, switches, and other switching elements that operate to direct data through the internetwork 110. For a more detailed description of the structure and operation of the Internet 110, one may refer to "The Internet Complete Reference," by Harley Hahn and Rick Stout, published by McGraw-Hill, 1994. However, one may practice the present invention on a wide variety of communication networks. For example, the internetwork 110 can include interactive television networks, telephone networks, wireless data transmission systems, two-way cable systems, customized computer networks, interactive kiosk networks, or ad hoc packet relay networks.

One embodiment of the internetwork 110 includes Internet service providers 108 offering dial-in service, such as Microsoft Network, America OnLine, Prodigy and CompuServe. It will be appreciated that the Internet service providers 108 may also include any computer system which can provide Internet access to a client 102. Of course, the Internet service providers 108 are optional, and in some cases, the clients 102 may have direct access to the Internet 110 through a dedicated DSL service, ISDN leased lines, T1 lines, digital satellite service, cable modem service, or any other high-speed connection to a network point-of-presence. Any of these high-speed services may also be offered through one of the Internet service providers 108.

In its present deployment as the Internet, the internetwork 110 consists of a worldwide computer network that communicates using protocols such as the well-defined Transmission Control Protocol ("TCP") and Internet Protocol ("IP") to provide transport and network services. Computer systems that are directly connected to the Internet 110 each have a unique IP address. The IP address

consists of four one-byte numbers (although a planned expansion to sixteen bytes is underway with IPv6). The four bytes of the IP address are commonly written out separated by periods such as "12.30.58.7". To simplify Internet addressing, the Domain Name System ("DNS") was created. The DNS allows users to access  
5 Internet resources with a simpler alphanumeric naming system. A DNS name consists of a series of alphanumeric names separated by periods. For example, the name "www.mdol.com" corresponds to a particular IP address. When a domain name is used, the computer accesses a DNS server to obtain the explicit four-byte IP address. It will be appreciated that other internetworks 110 may be used with the  
10 invention. For example, the internetwork 110 may be a wide-area network, a local area network, or corporate area network.

To further define the resources on the Internet 110, the Uniform Resource Locator system was created. A Uniform Resource Locator ("URL") is a descriptor that specifically defines a type of Internet resource along with its location. URLs  
15 have the following format:

*resource-type://domain.address/path-name*

where *resource-type* defines the type of Internet resource. Web documents are identified by the resource type "http" which indicates that the hypertext transfer protocol should be used to access the document. Other common resource types  
20 include "ftp" (file transmission protocol), "mailto" (send electronic mail), "file" (local file), and "telnet." The *domain.address* defines the domain name address of the computer that the resource is located on. Finally, the *path-name* defines a directory path within the file system of the server that identifies the resource. As used herein, the term "IP address" is intended to refer to the four-byte Internet  
25 Protocol address (or the sixteen-byte IPv6 address), and the term "Web address" is intended to refer to a domain name address, along with any resource identifier and path name appropriate to identify a particular Web resource. The term "address," when used alone, is intended to refer to either a Web address or an IP address.

In an exemplary embodiment, a browser, executing on one of the clients 102,  
30 retrieves a Web document at an address from one of the servers 104 via the

internet network 110, and displays the Web document on a viewing device, e.g., a screen. A user can retrieve and view the Web document by entering, or selecting a link to, a URL in the browser. The browser then sends an http request to the server 104 that has the Web document associated with the URL. The server 104 responds to the http request by sending the requested Web document to the client 102. The Web document is an HTTP object that includes plain text (ASCII) conforming to the HyperText Markup Language ("HTML"). Other markup languages are known and may be used on appropriately enabled browsers and servers, including the Dynamic HyperText Markup Language ("DHTML"), the Extensible Markup Language ("XML"), the Extensible Hypertext Markup Language ("XHTML"), and the Standard Generalized Markup Language ("SGML").

Each Web document may contain hyperlinks to other Web documents. The browser displays the Web document on the screen for the user and the hyperlinks to other Web documents are emphasized in some fashion such that the user can identify and select each hyperlink. To enhance functionality, a server 104 may execute programs associated with Web documents using programming or scripting languages, such as Perl, C, C++, or Java. A server 104 may also use server-side scripting languages such as ColdFusion from Allaire, Inc., or PHP. These programs and languages perform "back-end" functions such as order processing, database management, and content searching. A Web document may also include references to small client-side applications, or applets, that are transferred from the server 104 to the client 102 along with a Web document and executed locally by the client 102. Java is one popular example of a programming language used for applets. The text within a Web document may further include (non-displayed) scripts that are executable by an appropriately enabled browser, using a scripting language such as JavaScript or Visual Basic Script. Browsers may further be enhanced with a variety of helper applications to interpret various media including still image formats such as JPEG and GIF, document formats such as PS and PDF, motion picture formats such as AVI and MPEG, and sound formats such as MP3 and MIDI. These media formats, along with a growing variety of proprietary media formats, may be used to

enrich a user's interactive and audio-visual experience as each Web document is presented through the browser. The term "page" as used herein is intended to refer to the Web document described above, as well as any of the above-described functional or multimedia content associated with the Web document.

5           Figure 2 shows a block diagram of a server that may be used with the systems described herein. In this embodiment, the server 104 includes a presentation server 200, an application server 202, and a database server 204. The application server 202 is connected to the presentation server 200. The database server 204 is also connected to the presentation server 200 and the application server  
10   202, and is further connected to a database 206 embodied on a mass storage device. The presentation server 200 includes a connection to the internetwork 110. It will be appreciated that each of the servers may comprise more than one physical server, as required for capacity and redundancy, and it will be further appreciated that in some embodiments more than one of the above servers may be logical servers residing on  
15   the same physical device. It will further be appreciated that one or more of the servers may be at a remote location, and may communicate with the presentation server 200 through a local area or wide area network. The term "host," as used herein, is intended to refer to any combination of servers described above that include a presentation server 200 for providing access to pages by the clients 102.  
20   The term "site," as used herein, is intended to refer to a collection of pages sharing a common domain name address, or dynamically generated by a common host, or accessible through a common host (i.e., a particular page may be maintained on or generated by a remote server, but nonetheless be within a site).

          The presentation server 200 provides an interface for one or more  
25   connections to the internetwork 110, thus permitting more than one of the clients 102 (Fig. 1) to access the site at the same time. In one embodiment, the presentation server 200 comprises a plurality of enterprise servers, such as the ProLiant Cluster available from Compaq Computer Corp., or a cluster of E250's from Sun Microsystems running Solaris 2.7. Other suitable servers are known in the art and  
30   are described in Jamsa, Internet Programming, Jamsa Press (1995). The server

maintains one or more connections to the Internet 110, preferably provided by a tier one provider, i.e., one of the dozen or so national/international Internet backbones with cross-national links of T3 speeds or higher, such as MCI, UUNet, BBN Planet, and Digex. Each server may be, for example, an iPlanet Enterprise Server 4.0 from the Sun/Netscape Alliance. The presentation server 200 may also, for example, Microsoft's .NET technology, or use a Microsoft Windows operating system, with a "front end" written in Microsoft Active Server Page ("ASP"), or some other programming language or server software capable of integrating ActiveX controls, forms, Visual Basic Scripts, JavaScript, Macromedia Flash Technology multimedia, e-mail, and other functional and multimedia aspects of a page. Typically, the front end includes all text, graphics, and interactive objects within a page, along with templates used for dynamic page creation.

A client 102 (Fig. 1) accessing an address hosted by the presentation server 200 will receive a page from the presentation server 200 containing text, forms, scripts, active objects, hyperlinks, etc., which may be collectively viewed using a browser. Each page may consist of static content, i.e., an HTML text file and associated objects (\*.avi, \*.jpg, \*.gif, etc.) stored on the presentation server, and may include active content including applets, scripts, and objects such as check boxes, drop-down lists, and the like. A page may be dynamically created in response to a particular client 102 request, including appropriate queries to the database server 204 for particular types of data to be included in a responsive page. It will be appreciated that accessing a page is more complex in practice, and includes, for example, a DNS request from the client 102 to a DNS server, receipt of an IP address by the client 102, formation of a TCP connection with a port at the indicated IP address, transmission of a GET command to the presentation server 200, dynamic page generation (if required), transmission of an HTML object, fetching additional objects referenced by the HTML object, and so forth.

The application server 202 provides the "back-end" functionality of the Web site, and includes connections to the presentation server 200 and the database server 204. In one embodiment, the presentation server 200 comprises an enterprise server,

such as one available from Compaq Computer Corp., running the Microsoft Windows NT operating system, or a cluster of E250's from Sun Microsystems running Solaris 2.7. The back-end software may be implemented using pre-configured e-commerce software, such as that available from Pandesic, to provide  
5 back-end functionality including order processing, billing, inventory management, financial transactions, shipping instructions, and the like. The e-commerce software running on the application server 202 may include a software interface to the database server 204, as well as a software interface to the front end provided by the presentation server 200. The application server 200 may also use a Sun/Netscape  
10 Alliance Server 4.0. A payment transaction server may also be included to process payments at a Web site using third party services such as Datacash or WorldPay, or may process payments directly using payment server and banking software, along with a communication link to a bank. While the above describes one form of application server that may be used with the systems described herein, other  
15 configurations are possible, as will be described in further detail below.

The database server 204 may be an enterprise server, such as one available from Compaq Computer Corp., running the Microsoft Windows NT operating system or a cluster of E250's from Sun Microsystems running Solaris 2.7, along with software components for database management. Suitable databases are  
20 provided by, for example, Oracle, Sybase, and Informix. The database server 204 may also include one or more databases 206, typically embodied in a mass-storage device. The databases 206 may include, for example, user interfaces, search results, search query structures, lexicons, user information, and the templates used by the presentation server to dynamically generate pages. It will be appreciated that the  
25 databases 206 may also include structured or unstructured data, as well as storage space, for use by the presentation server 200 and the application server 202. In operation, the database management software running on the database server 204 receives properly formatted requests from the presentation server 200, or the application server 202. In response, the database management software reads data  
30 from, or writes data to, the databases 206, and generates responsive messages to the



requesting server. The database server 204 may also include a File Transfer Protocol (“FTP”) or a Secure Shell (“SSH”) server for providing downloadable files.

While the three tier architecture described above is one conventional architecture that may be used with the systems described herein, it will be appreciated that other architectures for providing data and processing through a network are known and may be used in addition to, or in conjunction with, or in place of the described architecture. Any such system may be used, provided that it can support aspects of the image processing system described herein.

Figure 3 shows a page that may be used as a user interface. The page 300 may include a header 302, a sidebar 304, a footer 306 and a main section 308, all of which may be displayed at a client 102 using a browser. The header 302 may include, for example, one or more banner advertisements and a title of the page. The sidebar 304 may include a menu of choices for a user at the client 102. The footer 306 may include another banner advertisement, and/or information concerning the site such as a “help” or “webmaster” contact, copyright information, disclaimers, a privacy statement, etc. The main section 308 may include content for viewing by the user. The main section 308 may also include, for example, tools for electronically mailing the page to an electronic mail (“e-mail”) account, searching content at the site, and so forth. It will be appreciated that the description above is generic, and may be varied according to where a client 102 is within a Web site related to the page, as well as according to any available information about the client 102 (such as display size, media capabilities, etc.) or the user.

A Web site including the page 300 may use cookies to track users and user information. In particular, a client 102 accessing the site may be accessed to detect whether the client 102 has previously accessed the page or the site. If the client 102 has accessed the site, then some predetermined content may be presented to the client 102. If the client 102 does not include a cookie indicating that the client 102 has visited the site, then the client 102 may be directed to a registration page where information may be gathered to create a user profile. The client 102 may also be

presented with a login page, so that a pre-existing user on a new client 102 may nonetheless bypass the registration page.

The site may provide options to the client 102. For example, the site may provide a search tool by which the client 102 may search for content within the site, or content external to the site but accessible through the internet network 110. The site may include news items topical to the site. Banner ads may be provided in the page 300, and the ads may be personalized to a client 102 if a profile exists for that client 102. The banner ads may also track redirection. That is, when a client 102 selects a banner ad, the link and the banner ad may be captured and stored in a database. The site may provide a user profile update tool by which the client 102 may make alterations to a user profile.

It will be appreciated that the foregoing description has been generic. A user interface for a medical image processing system will now be described in more detail. It will be appreciated that the interface may be embodied in any software and/or hardware client operating on a client device, including a browser along with any suitable plug-ins, a Java applet, a Java application, a C or C++ application, or any other application or group of applications operating on a client device. In one embodiment, the user interface may be deployed through a Web browser. In one embodiment, the user interface may be deployed as an application running on a client device, with suitable software and/or hardware for access to an internet network. In these and other embodiments, certain image processing functions may be distributed in any suitable manner between a client device and one or more servers, as will be explained in further detail below.

In the example embodiment described below, each workspace, including a patient workspace, an atlas workspace, a reference workspace, and a results workspace, may be selected using tabs such as those provided for Windows applications. Each workspace will now be explained in further detail.

Figure 4 shows a patient workspace of a user interface. As described above, the interface 400 may be deployed as an application running on a local machine, as a remote service run from an application service provider, as a Web-based resource

accessible over an internetwork, or any other mode suitable for use at a client device. Functionality may be distributed in any suitable manner between the client device and one or more remote resources such as databases, servers, and the like. It will be appreciated that the interface 400 depicted in Fig. 4 is an example, and that  
5 other arrangements of the interface may be used consistent with the systems described herein. It will also be appreciated that menus, submenus, and other interface screens may be usefully employed to support the functionality of the interface, such as menus for controlling scoring and relevancy of search results.

10 A patient workspace 402 may include a worklist area 404 and an image display area 406, along with any suitable navigation aids, such as the vertical scroll bar depicted on the right hand side of the image display area 406. The worklist area 404 may include a button for accessing a worklist managing interface through which a user may add, remove, or otherwise manage and dispose of cases, including  
15 providing diagnostic conclusions and any other findings or observations. A worklist displayed in the worklist area may include one or more cases for review, such as neural radiology MR studies. Once a case has been reviewed and, for example, a diagnosis reached, the case may be dispatched from the worklist using the button.

The image display area 406 may display one or more images, such as a series  
20 of images in an MR study. Through the interface 400, a user may select on of the series of images for enlarged display. In this manner, a user may navigate through images such as provided in an MR study, and review in detail areas of possible interest. In one embodiment, a user may select one or more regions of interest graphically within the image display area 406. Regions of interest, such as possible  
25 pathologies or other abnormalities, may be demarcated as points, predetermined geometric shapes (e.g., squares, triangles, ellipses, etc.), or as hand-drawn contours.

Fig. 5 shows an atlas workspace of a user interface. The atlas workspace 502 may include a view selection menu 504 and an atlas view space 506. The view selection menu 504 may be used to specify, for example, an anatomical region, an

imaging modality, a pathology, or any other criteria useful for selecting a library of images.

The atlas view space 506 may provide views of subject matter selected in the view selection menu 504. For example, the atlas view space 506 may display an axial view, a coronal view, a sagittal view, and a three-dimensional cut-away view based upon images such as MR images. Each view within the atlas view space 506 may include one or more navigation tools. For example, the axial view, the coronal view, and the sagittal views provided as examples in Fig. 5 may include a scroll bar for navigating through different slices of an MR study. In a three-dimensional cut-away view, also shown as an example view in Fig. 5, each plane of the cut-away may be determined by one of the other views in the atlas view space 506. The three-dimensional view, or tri-plane view, may also be separately controlled to display planes of varying depth in each dimension. The three-dimensional view may be rotated about its axes using, for example, keyboard input from the client device or mouse-over movements within the three-dimensional view space. Labels may also be associated with regions within the atlas. Labels may then be automatically displayed within the views based upon mouse positioning within the view. Or a user may initiate a query for a label corresponding to a particular location.

Figure 6 shows a reference workspace of a user interface. The reference workspace 602 may include a text area 604 and an image area 606. The text area 604 may include, for example descriptive, educational, or diagnostic information for normal anatomy and abnormalities. Text within the text area 604 may be hyperlinked to other reference materials locally maintained for access through the reference workspace 602, or remotely accessible through the internetwork. The text area 604 may be accompanied by one or more buttons such as a back button for moving backward through the text, a forward button for moving forward through the text, and an index button for accessing and index or table of contents for reference materials available through the reference workspace 602. One or more scroll bars may also be provided for manually navigating through a section of reference text.

The image area 606 may display a series or a single image relating to text displayed in the text area 604. Navigational tools may be provided for user control of images displayed in the image area 606, such as a scroll bar for navigating through slices of an MR study. A matching button may be provided for matching a  
5 current image, e.g., an image accompanying the reference text, to other images and/or clinical data stored within the system.

Figure 7 shows a results workspace of a user interface. The results workspace 702 may include a query area 704, a results area 706, and an image area 708. The query area 704 may show one or more images for analysis through the  
10 user interface 400. The results area 706 may show one or more matches to the images, and/or associated clinical data, from the query area 704. The matches may be organized, such as by pathology or by queried database, and may be ranked according to a score derived from matching criteria. The image area 708 may display a query image associated with a subject listed in the query area 704,  
15 including any regions of interest identified within the image(s). The image area 708 may also display a result image, including one or more images and associated clinical data matched to the query image, such as through the matching techniques described below. The image area 708 may also display thumbnails of one or more result images along with any descriptive information, such as an associated  
20 pathology, a similarity ranking, a relative matching score, clinical data, and biographical data for an associated patient. A user may select one of the thumbnails for display as a result image.

In one aspect, the user interface 400 provides a platform for multi-modal matching or a multi-modal search engine. Each mode may provide a type of  
25 matching against a database of images and other information, which may be, for example, an image database containing images pre-processed for matching, as described in more detail below. Atlas matching may be provided, such as in the atlas workspace 502, in which images may be retrieved that match the position (in one, two, or three reference planes) being viewed in, or selected within, a three-  
30 dimensional object displayed within the atlas. Similarity matching may be provided,

such as in the reference workspace 602, in which images may be retrieved that match the appearance of an image being viewed. Similarity may in this context be measured using matching and scoring techniques described below, or any other technique for evaluating similarity between images and image data. Diagnostic  
5 matching may be provided, such as in the results workspace 702, in which a diagnosis is generated for a query image. The diagnosis may include similar images retrieved from the database, and further include more than one diagnosis, scored according to comparison with images and diagnoses available in the database

It will be appreciated that the above user interface may be used as a platform  
10 for workflow management in a clinical or other medical setting, as described below in more detail in reference to Fig. 11.

Having described an interface for using an image matching processing system in a medical context, a system for matching images is now described in further detail. The following matching techniques may be used in combination with  
15 the user interface described above in order to provide, for example, matched images in the results workspace 702.

Figure 8 is a flow chart showing a process 800 for processing images according to the systems described herein. The following example embodiment describes a method for processing neural radiology images obtained through  
20 magnetic resonance imaging. However, it will be appreciated that the techniques described herein may apply to a broad range of anatomical images obtained through a number of different imaging modalities, including x-ray images, computed tomography images, magnetic resonance imaging, ultrasound, and so forth. Any of these images, as well as non-medical images and other complex data or data  
25 structures, may be processed using the systems described herein.

In the systems described herein, there are periodic references to an x-direction, a y-direction, or a z-direction, along with related references to an x-position, a y-position, or a z-position and mathematically derived values such as an x-gradient or a y-gradient. It will be appreciated that an x-direction and a y-  
30 direction generally refer to two orthogonal axes in a planar image, and that a z-

direction refers to a third axes perpendicular to the planar image. Other coordinate systems may be used with the systems described herein, such as the polar coordinates used for registration, which lie in the planar image.

It will be appreciated that, while the following description refers generally to cerebral magnetic resonance images, other image types are possible. For example, computerized tomography images may be used. Additionally, images may be taken from various anatomical regions, such as neck images, spine images, or musculo-skeletal images. All such image types are intended to fall within the scope of the systems described herein. Furthermore, the systems described herein may be extrapolated to full three-dimensional figures, as distinguished from the series of planar two-dimensional images typical of an MRI study.

As shown in step 802, database images may be provided. These may be, for example, an axial MR study including digitized images of each slice of the study. Additional data may be associated with each image, such as imaging details (e.g., pulse sequences such as T1, T2, FLAIR, and PD, or orientations such as axial, sagittal, or coronal), anatomy (e.g., brain, torso, arm, chest, etc.), patient data (e.g., age, height, weight, patient identification, gender, diagnosis, clinical reports, etc.), and any other data (e.g., physician, number of images, date, time) that may relate to the images. The additional data may be provided with each study, as in DICOM headers used for medical images, or may be added (manually or automatically) as supplemental information for each study.

As shown in step 804, each image may be pre-processed. This may generally include normalization, segmentation, and feature extraction, each of which will be discussed in more detail below.

Normalization may begin with mask extraction using any suitable boundary detection technique to obtain a mask of the outside perimeter of an imaged anatomical region. Within the mask area, grayscale intensity may be normalized for enclosed pixels. The mask image may then be rotated and registered to a global coordinate system. This may be, for example, a two-dimensional registration for each slice of an MR study, applying a non-rigid, global transformation, to an atlas

derived from sample data, or to some predetermined geometry approximately descriptive of the images. One suitable geometry for cortical images is an ellipse. A number of mathematical techniques are known for performing non-rigid registration, including deformable model based registration, geometrical or other landmark based registration, voxel property based registration, and so forth. Transformations  
5 between registered and unregistered images may be affine, projective, or elastic. Any such technique may be used with the systems described herein, provided that (visual) features within different images may be adequately superimposed in the global coordinate system to allow discrimination of pathological variations, relative  
10 z-position, and any other information that might be usefully extracted from images and further processed using the other techniques described herein.

Once an image has been registered to a global coordinate system, it may be sampled in discrete sections within the transformed image space. In MR imaging, for example, an approximately polar coordinate system (e.g., elliptical) may be used  
15 for the global coordinate system, and tiles may be arranged for discrete sampling along axes of the global coordinate system. Arrangement of tiles within a polar, elliptical, or nearly elliptical coordinate system may be useful for MR images. In this manner, tiles may be arranged to coincide with regions of interest known to be associated with pathologies.

Referring briefly to Fig. 9, several possible arrangements of tiles for an MR image are shown. A slice of an MR study 902 may be masked, and the mask may be registered using a non-rigid transformation to an approximately elliptical shape. It will be appreciated that the transformed dimensions may be arbitrary, particularly where it is not used for visualization. However, suitable selection of a coordinate  
25 system for the transform may reduce computational costs, simplify registration, and map more effectively to the structure of the imaged subject, and implicitly, likely areas of diagnostic interest. One suitable transform is to the approximately elliptical shape shown in Fig. 9. The shape may then be tiled in any suitable manner. For example, a relatively coarse tiling 904 may use forty-nine tiles to cover the masked,  
30 registered image. Finer tiling 906, 908, 910 may be used where regions of interest



have smaller dimensions. It will also be noted that tiles may be strategically arranged to coincide with likely regions of interest, such as to cover an area directly north (as oriented in Fig. 9) from the center, which corresponds to an indication of Pituitary Microadenoma at the depicted z-position.

5 Referring again to Fig. 8, image data in each tile may be sampled, and feature vectors may be extracted, as shown in step 805. Sampling may be, for example, an  $N \times N$  matrix of grayscale values from within each tile. Feature vectors may include, for example, tile size, mean signal intensity (as normalized), standard deviation of signal intensity, mean edge magnitude, fraction of edge points above a  
10 threshold, mean x-direction gradient, mean y-direction gradient, mean absolute value of surface curvature, and mean absolute value of levelset curvature. In one embodiment, forty-nine tiles are used to cover the mask space, and ten feature vectors are calculated for each tile.

As shown in step 806, the database images may be segmented by, for  
15 example, identifying regions of interest within each image. This may be performed manually through, for example, a graphical user interface. Manual entry may be particularly useful where there is no *a priori* information about the database images and expected regions of interest. As will be described in further detail below, this may also be automated by modeling images with known regions of interest, and  
20 applying these models to new images to test for the presence of certain regions of interest. Regions of interest may be designated as points, as geometric shapes (e.g., square, triangle, circle, ellipse, trapezoid, etc.), or as hand-drawn or otherwise free-form open or closed curves.

As shown in step 808, the database images may be inspected to determine a  
25 relative z-position of each slice, such as of an MR study. This may be performed manually through, for example, a graphical user interface. Manual entry may be particularly useful where there is no *a priori* information about the database images and expected relative z-positions. As will be described in further detail below, this may also be automated by modeling images with known relative z-positions, and  
30 applying these models to new images to determine a relative z-position.

More generally, while steps 806 and 808 depict labels of images for regions of interest and relative z-position, any labels may be similarly added to images, or regions of images, to correspond to observed characteristics (e.g., diagnosis, diagnostic significance of a region of interest, unusual characteristics, clinical indications associated with a patient, and so forth) or other global attributes of the images. These labels may be added manually to a database by inspecting images and associated data, and adding any suitable or desired labels. For example, a qualified physician may examine images and other patient data, and label images or locations therein as indicative of a specific pathology. Models may also be constructed from manually labeled images to automate subsequent labeling of data. It will also be appreciated that any known computational techniques may be applied for automated labeling, without regard to computational cost, without affecting subsequent matching and modeling steps described herein.

As depicted in step 810, a filter may be provided for one or more databases of images that have been pre-processed as described above. The filter may receive data relating to the query image, as described in more detail below, and may prepare a subset of data from the database of images that is to be used for subsequent matching. For example, the database of images may be filtered based upon pathologies, orientation and sequence, z-location range, and region of interest. Filtering by pathology may exclude all normal image cases (i.e., those not exhibiting any pathology), or may exclude all images not exhibit a specified pathology where, for example, a particular pathology is being tested for. Orientation and sequence filtering may remove images obtained, for example, on a different axis from the query image, or using a different imaging modality. For example, MR imaging may employ several different pulse types, sequences, and so forth, and each type may produce visually different results for the same region, thus rendering image-based comparisons meaningless, or at least less useful. Filtering may also be performed by z-location using, for example, the derived or manually added z-position data provided in step 808. Filtering may also be performed by region of interest, using, for example, the derived or manually added region of interest data provided in step

806. Using these filters, subsequent matching steps may be performed on a subset of the complete database of images.

As shown in step 812, a query image may be provided. This may be, for example, an axial MR study including digitized images of one or more slice of a study. Additional data may be associated with each query image, such as imaging details (e.g., pulse sequences such as T1, T2, FLAIR, and PD, or orientations such as axial, sagittal, or coronal), anatomy (e.g., brain, torso, arm, chest, etc.), patient data (e.g., age, height, weight, patient identification, gender, diagnosis, etc.), and any other data (e.g., physician, number of images, date, time) that may relate to the images. The additional data may be provided with each query image, as in DICOM headers used for medical images, or may be added (manually or automatically) as supplemental information for each query image.

As shown in step 814, each query image may be pre-processed. This may generally include normalization, segmentation, and feature extraction, each is discussed above with reference to database images. Further steps such as feature extraction shown in step 816, region of interest determination 818, relative z-position 820, and any other labeling with observed characteristics or other global attributes may be performed on the query image(s). Regions of interest, for example, may be identified by a user with a single point-and-click mouse operation at a client device where the query image is being reviewed. It will be appreciated that, in a client/server architecture such as those discussed above, a browser-based or application-based client may perform steps 814-820 at a client device, with matching performed at a remote location accessible through the internetwork. Matched images may then be returned to the client device for review by a user at that location. Other arrangements are possible such as receiving unprocessed query images at a workstation that has the database of images, filter, and so forth locally available. Where the filtered database becomes sufficiently small, the output of the filter may also be provided to a client device for local matching at the client device. Full images may then be retrieved for review at the client device on an ad hoc basis,

depending on results of a match performed at the client device. All such variations are intended to fall within the scope of the systems described herein.

As shown in step 822, matching may be performed between the query image and one or more images from the database of images. More particularly, a pre-processed query image and accompanying labels may be compared to a filtered set of pre-processed images from the database of images. As noted above, the labels for the query image may be used, in part, to filter data from the database of images. In one embodiment, a matching algorithm may receive from the filtering step 810 any features extracted for images that satisfy filter constraints. These feature vectors may be compared to one or more corresponding feature vectors from the query image. A scoring mechanism for this process may be chosen from a wide spectrum of procedures, such as a simple two-norm, inner product, or correlation, or other similarity measure. Elements of each feature vector may be weighted according to an empirical or formulaic evaluation of results. Results of the scoring process may be further processed, such as by removing results below a certain matching score threshold, and by sorting according to score, pathology, and so forth.

In certain cases, the matching may be improved by appropriately weighting one or more feature vectors. For example, a Mahalanobis distance may be applied to normalize weights for feature vectors used in the model. In application, matching may be performed by computing a weighted norm of a difference between feature vectors for a query image and each image that the query image is tested against. The weighted norm may be a Mahalanobis distance derived from a covariance matrix of database feature vectors. More generally, the weighted norm may be any function of one or more labels or other data, including spatial position, pathology, and so forth.

As shown in step 824, results may be evaluated. This may be conducted by a user through, for example, the user interface described above with reference to Figs. 4-7. Once a match has been made, any or all of the data relating to the matched image may be retrieved and reviewed by a user, including a full MR study, associated clinical data, biographical patient data, labels, and so forth. The evaluation may include references to an atlas, reference materials, and other data

that may be available through the user interface described above. Furthermore, any conclusions or other findings may be associated with the query image and stored along with the query image in a suitable database.

At the conclusion of the evaluation, findings may be provided to a diagnosis  
5 model as shown in step 826. For example, a neural radiologist may review a query image, along with all associated clinical and other data. The radiologist may further review matched images and associated data obtained through the matching process of step 822. If the radiologist concludes that a particular diagnosis is appropriate, this indication may be provided to a modeling system along with the query image,  
10 associated data, and any other data derived therefrom, such as feature vectors, regions of interest, relative z-positions, and any other labels or other attributes. Query images stored in this fashion may provide, for example, a ground-truth database of diagnosed cases from which diagnostic models may be obtained using techniques described below. A diagnostic model may be used to automate diagnosis  
15 of subsequent query images, or to recommend one or more diagnoses to a qualified physician.

As shown in step 828 other modeling may be performed. This modeling may be performed on any combination of image data from the database of images, pre-processed image data, including feature vectors, regions of interest, relative z-  
20 position, and any other labels, as well as other data associated with the images as described above. Modeling may generally include training a model to associate one or more inputs with one or more outputs. For example, feature vectors may be associated with a label, such as relative z-position, or with a diagnosis of a pathology. Techniques for modeling include, for example, regression analysis with  
25 regression coefficients determined through least squares or partial least squares, neural networks, fuzzy logic, and so forth. Regression modeling, for example, has been usefully applied to feature vectors for MR images to identify z-location, to note the presence of contrast enhancing agents, and to discriminate between T1 and MRA images. These models may be applied, for example, to automate labeling processes  
30 and to label new databases, or to check the accuracy of labels for databases that

already include the label information. Modeling is discussed in further detail with reference to Fig. 10.

It will be appreciated that the above process 800, may be realized in hardware, software, or some combination of these. The process 800 may be realized  
5 in one or more microprocessors, microcontrollers, embedded microcontrollers, programmable digital signal processors or other programmable device, along with internal and/or external memory such as read-only memory, programmable read-only memory, electronically erasable programmable read-only memory, random access memory, dynamic random access memory, double data rate random access  
10 memory, Rambus direct random access memory, flash memory, or any other volatile or non-volatile memory for storing program instructions, program data, and program output or other intermediate or final results. The process 500 and the shape processor 200 may also, or instead, include an application specific integrated circuit, a programmable gate array, programmable array logic, or any other device that may  
15 be configured to process electronic signals.

Any combination of the above circuits and components, whether packaged discretely, as a chip, as a chipset, or as a die, may be suitably adapted to use with the systems described herein. It will further be appreciated that the above process 800 may be realized as computer executable code created using a structured  
20 programming language such as C, an object-oriented programming language such as C++ or Java, or any other high-level or low-level programming language that may be compiled or interpreted to run on one of the above devices, as well as heterogeneous combinations of processors, processor architectures, or combinations of different hardware and software. The process 800 may be deployed using  
25 software technologies or development environments including a mix of software languages, such as Microsoft IIS, Active Server Pages, Java, C++, Oracle databases, SQL, and so forth. In addition, it will be appreciated that, as where noted above, certain steps to the process 800 may be realized in part through input from a human user, such as where initialization of a model is to be based on diagnoses from a  
30 trained radiologist.

Figure 10 is a flowchart of a process for organizing databases. While the following example details a technique for organizing cerebral MRI images using statistical techniques to develop regression coefficients, it will be appreciated that the general approach described herein may have general application to organization and searching of large databases, and in particular to databases of medical images. For example, a model that relates feature vectors to a relative z-position may be used to automate labeling of relative z-positions for slices of an MR study. Or, a model that relates regions of interest and feature vectors to a diagnosis may be used as a diagnostic tool. These and other applications are intended to fall within the scope of the system described herein. As an example, text-base modeling may be performed, in which textual patient reports may be analyzed and modeled for labeling by diagnosis or other criteria.

A source database 1010 may be used to obtain derived data, as shown in step 1020. The source database 1010 may be for example, a database of MR studies, along with clinical and other patient data. Derived data may be data obtained through any predetermined function. For example, the feature vector extraction described above with reference to Fig. 8 produces derived data in the form of feature vectors for an image, or for a region of an image. The derived data may optionally be stored in the source database 1010 along with the images.

As shown in step 1030, images may then be labeled. This may include manual labeling of regions of interest or z-positions, as described above. This may also include identification of data already associated with images in the source database 1010, such as a diagnosis, or patient data such as age, sex, height, weight, and so forth.

As shown in step 1040, labels obtained in step 1030 and derived data obtained in step 1020 may be applied to a model, such as a regression model or a neural network model, and coefficients may be obtained that relate inputs (e.g., derived data) to outputs (e.g., labels). The model may employ, for example, linear regression, or any other statistical learning methodology. The resulting trained model may be applied to a target database 1050 as described below. The model may

be adapted to different types of outputs. For example, a model with a scalar output may be used for a one-dimensional result, such as z-position. Such a model may also be applied to binary inquiries. For example, where a model is constructed to test for the presence of a contrast agent, a larger value may correspond to a contrast agent while a smaller value may correspond to no contrast agent. Where more information is required, such as an x-location and y-location (or angle and distance) for a region of interests, a multi-dimensional output may be appropriate.

As shown in step 1060, data may be derived from the target database 1050. While any technique may be used to derive data from the target database 1050 in step 1060, as well as to derive data from the source database 1010 in step 1020, the same technique should be used in both cases so that the model trained in step 1040 will yield meaningful results. As shown in step 1070, the model may be applied to the data derived from the target database 1050, and the output of the model may be used to label data from which the derived data was obtained. This procedure may be useful, for example to prepare a new database of MR images for use with the image processing systems described above. More particularly, operations such as manual data entry of relative z-positions for new images may be replaced with automated determination of relative z-position, and images automatically labeled with z-positions in this manner may be added to a source database for image matching.

A model trained as described above may also, or instead, be used to test the integrity of pre-existing labels in databases. Following the example above, a database may already include relative z-positions for images in the database. The model may be applied as above to independently determine relative z-positions. The relative z-positions may then be compared in a number of ways. For example, the model output may be used to replace pre-existing labels, or may be used to provide labels for any images missing this information. The results may be reported statistically as deviations from expected results. Or the results may be used to exclude from subsequent searches those images with pre-existing labels significantly different from a corresponding model output. The significance of the difference



may be determined using an absolute or relative threshold for excluding certain database records.

The above approach to modeling databases may be applied in a data mining system. For example, a physician may search image databases for images matching  
5 a particular patient by applying the modeling techniques above. The physician may, at the same time search text references, clinical histories, and any other data from other data sources, and gather search results into a library of possibly relevant materials. Further, a number of different models may be used to assist in searching  
10 databases having different data types. Image data may include, for example, neural CT scans, neural MR studies, and other image types. Modeling may also be applicable to other data types, such as text-based patient histories, provided that the data can be labeled with one or more observed characteristics, and processed into some type of derivative data for which a model can be accurately trained to generate the labels. There is accordingly described herein a data mining system in which one  
15 or more target databases can be modeled and queried, either alone or in conjunction with other searches.

Figure 11 shows a state diagram for a workflow management system. The workflow management system 1100 may be formed around several states of a user interface, such as a patient state 1102, an atlas state 1104, a reference state 1106, and  
20 a results state 1108. As indicated in Fig. 11 by arrows interconnecting these states, a user may navigate between states in any suitable manner while working to conclude with a diagnosis for a patient.

For example, the system may be initiated in the patient state 1102, where a radiologist may receive an MR study for review in a user interface, such as the  
25 patient workspace 402 of Fig. 4. The radiologist may navigate to the atlas state 1104, where a workspace such as the atlas workspace 502 of Fig. 5 may be used as a guide to visual review of an anatomical area imaged by the MR study. The radiologist may navigate to the reference state 1106, where a workspace such as the reference workspace 602 of Fig. 6 may be used by the radiologist to investigate any

potential diagnoses. Each state may be selected by, for example, selecting a window tab in a Windows environment, or using a navigation panel in a Web browser.

The results state 1108, which may be presented in a workspace such as the results workspace 702 of Fig. 7, may be reviewed by the radiologist, or other  
5 clinician, as a further diagnostic aid. The results state 1108 may be reached using navigation methods noted above, where image matches and associated data, such as clinical data, from a previous search may be reviewed. The results state 1108 may also be reached through an explicit search instruction provided in one of the other states, e.g., the patient state 1102, the atlas state 1104, or the reference state 1106. A  
10 search from the patient state 1102 may be for images in an image database that match image data for a patient currently under review, and may include, for example, a user-specified region of interest. A search from the atlas state 1104 may be for all images corresponding to a user-specified location within an atlas displayed in the atlas workspace 502. A search from the reference state 1106 may be a search  
15 upon a user provided criterion, such as a pathology, a location, and so forth.

A radiologist using the system for diagnosis may traverse the states as desired, moving for example, from the results state 1108 to the reference state 1106 for further information about one or more pathologies, until a conclusion is reached concerning a current MR study. The MR study may then be dispatched from the  
20 system, along with any findings, and additional MR studies may be retrieved for analysis. Where MR studies are provided directly in digital form, they may be pre-processed as described below and provided to the workflow management platform for review by a clinician in a fully paper-less system.

It will be appreciated that other applications for the above systems are  
25 possible. For example, where all images are registered to a global coordinate system as described above, a spatial probability map may be created for a pathology. That is, each region of interest in an image associated with the pathology may be combined using known techniques, and spatial distributions of the regions of interest combined so that the likelihood of a region of interest (such as a lesion) appearing in  
30 a certain location for a certain pathology may be determined. Following this

approach through the entire coordinate system, a complete spatial probability map may be derived for a particular pathology. The map would reveal, for each location within the coordinate system, the probability of a lesion appearing for the pathology. Such a map may be used as a diagnostic aid, or as an aid to identifying regions of interest in manual or automated labeling systems such as those described above.

It will further be appreciated that a number of different architectures may be based upon the system described herein. For example, a client providing the user interface and query image pre-processing described herein may be licensed for remote use through a Web browser plug-in, remote use through an application service provider, or remote use through a proprietary local client. The software and/or hardware system may, instead, be sold or licensed for use in its entirety as a clinical tool for use on a local or corporate area network of a hospital.

The systems described herein may be extended to operate as a just-in-time search engine. One of the most valuable resources in a medical care environment is time. A typical database system usually involves the generation and submission of a user-defined query that is processed by the database while the user waits for the results. This commonly used model is appropriate in situations where the query and the results are relatively small in size to transmit, and the processing of the query occurs quickly. However, in the medical domain, there are numerous delays that can occur between the time when the query is initially submitted and the time when the results have been received.

The query consists, in part, of the results of patient tests, and may include all forms of digital information. Unlike text-only queries that are relatively small in size, three-dimensional medical imagery is slow to transmit over existing networks, and can be a major source of delay. While the query generally consists of a small number of images, the database may return a large number (20 or more) of relevant images as results, causing significant delay in receipt of the entire reference set. The database may also return other information, such as three-dimensional models of the anatomical structures present in the reference images, or case-independent information mined from online medical encyclopedias. Together, all of the

multimedia information, such as text, imagery, audio, video, contained in the returned knowledge base can account for a substantial amount of data and thus requires time to transfer over the networks.

In addition to limitations in the network bandwidth, the analysis of the query  
5 also contributes to the delay. In order to determine which subset of the database is relevant to the patient's condition, portions of the query are matched to the elements of the database to find similarities. For example, an undiagnosed magnetic resonance (MR) image of a patient may be compared to MR images in the database that are tagged with diagnoses. A certain diagnosis of the query image is greatly  
10 supported when multiple images from the database with the same diagnosis all show a high degree of similarity to the query image. One practical challenge in this use of this matching system lies in the fact that the analysis time may not be constant for all cases. Pathology present in some imagery may be rapidly identifiable, while other cases may require more resources to perform sufficient analysis. When longer  
15 analysis time is required, it is unreasonable for the user to be idle, waiting for the results.

While many database-query systems consist of a one-stage query-response paradigm, a more general, and potentially more powerful approach involves a hierarchical, multi-stage information gathering protocol. For example, when an  
20 image matching system generates a list of similar medical images with corresponding diagnoses, the database system may in turn query other available resources based on the matched diagnoses and consolidate multimedia information to create a unified presentation of the results. While this multistage information mining process is a useful diagnostic aid, it may require sufficient time that if the  
25 query is submitted at the time of evaluation, the physician may not be able to wait for the results.

The delay inherent to remote medical information gathering may seem too large for practical use by physicians. However, a mining framework other than the standard method of submitting the query and waiting for the result provides a  
30 solution. In practice, there is generally a delay between the time in which medical

tests are performed and when a physician evaluates the results. For example, a radiologist technician may acquire a set of medical images of a patient hours before the radiologist analyzes the imagery. If the medical image is sent as a query to the database system at the time of acquisition, the delay between acquisition and  
5 evaluation allows for ample time for the query image to be sent to the remote database, for identification of relevant reference images from the database, and for transmission of the reference images back to the remote location.

Generally, the amount of time between the receipt and evaluation of test results varies from case to case. In non-emergency situations, the physician  
10 generally does not review the patient's medical imagery or test results as soon as they become available. Therefore, the existing delay between the acquisition and medical evaluation can be used effectively for transmission of data and extensive processing of the query, as well as for scheduling the transmission with other, more urgent requests.

15 In one embodiment of the present invention, the mining system consists of a gateway medical station that collects key input data for a patient from multiple sites within the medical facility, immediately upon availability of the data. The gateway also plans, schedules, and initiates all remote queries to most effectively use the resources available and to balance the needs of the individuals who submitted the  
20 queries. The key input data collected by the gateway usually consists of medical test results in the form of digitally encoded textual and/or graphical information. For example, medical images, such as magnetic resonance (MR) or computed tomography (CT) of a patient are acquired and stored digitally in the radiology department, while the patient's blood tests results are stored on the blood lab  
25 computer system. The gateway medical station would collect this and other patient-specific digital information as it is generated.

In addition to collecting the patient-specific information, the gateway also processes and organizes it into a set of database queries. The queries are then submitted to the relevant remote medical databases in order to retrieve additional  
30 useful resources for presentation to the physician at the time of evaluation. In

essence, the “*just in time*” system uses the available medical information to anticipate various retrieval requests that may be submitted by the physician at a later time, during the case evaluation. The choice of what information to mine depends on information contained in the patient’s records and the types of medical images acquired and tests performed. The technicians running the various diagnostic procedures can choose from a set of *a priori* mining procedures and can also submit requests for specific information. By processing, analyzing, and fetching relevant information immediately, the system leverages the elapsed time that generally occurs between the acquisition of the medical tests and evaluation of the results.

The scheduling algorithm incorporated into the data mining system balances the needs of the users with the available resources in a variety of ways. Given the potentially large number of simultaneous requests and large amount of data being transferred, the requests cannot simply be processed in one chunk as they arrive, since the urgency of information requests will vary significantly from case to case. In addition to test results, the gateway medical station also receives as input the expected time that the results will need to be reviewed by the physician. These parameters of the mining process are set at the time of acquisition either by the technician or doctor, in order to encode the urgency of the request. With this information, the gateway prioritizes the queries and schedules the submission of the queries, taking into account the expected length of time for transmission and processing of the query, as well as for receipt of the results. The information from any one patient may also be transformed into multiple queries submitted to multiple databases or preempt more urgent queries, while still ensuring receipt of the necessary information in the allotted amount of time. The scheduling system also takes into account the fact that information received from one database may be used as part of a future query to another database, requiring ample time to perform the multi-stage data mining.

In order to transmit information in the correct form to the appropriate database systems, the gateway stores information about the various remote sites, such as the type of data in the remote database, the types of queries allowed, and the

service level agreement (SLA) between the local and remote institutions. The SLA is a contract that determines the cost, response time, and resource constraints when a query is sent to a given search engine. For a standard type of query, the remote site defines the cost and how long before the results are returned. In urgent situations, the medical facility may need a response in less time than guaranteed by the basic terms of the SLA. Often the SLA includes the option of a rush order, at increased expense, guaranteeing the result in less time. This is most practical when the physician or radiologist includes with the data a constraint that the results are needed by a certain time (*e.g.* before a scheduled surgery). In general, the SLA provides a list of response times as a function of query types and expense that the gateway processes and uses as part of the prioritizing and scheduling algorithm to mine the needed remote information. The gateway scheduling system needs to take into account not only the information provided in the SLA, but also the policies as defined by the medical facility for transmission of such queries. For example, the medical facility may limit the expenditure on certain types of queries, so as to not have every request marked as urgent.

During times when there is a large load on the mining coordination system and external networks or when available time is severely limited, the gateway system will prioritize which data and query types are most relevant to transmit given the finite resources. This process includes organization of the data in a way to most effectively satisfy all the supplied constraints while accruing the least cost from the remote servers. Organization of the data involves stratification of the various requests, and in some cases, breaking apart the data into packets, and interleaving the packets as necessary to stream the information to allow remote processing to begin as soon as possible. This data organizing stem is also important since a remote site may limit the amount of data transmitted to its site in a given amount of time so as not to overload its own network pipeline, as specified in the SLA.

Once all of the query results for a given case have been received by the gateway mining system, the gateway then dispatches the information back to the appropriate local medical data repositories. The mined information will then be

available when the physician is prepared to analyze the patient's medical data for diagnosis and treatment purposes. In contrast with prior art, this method provides timely access to practically unlimited number of medical knowledge bases and, at the same time, it provides the characteristics of real time medical information mining services.

The systems described herein may also be applied to create self-forming groups based upon common interests, such as medical conditions. Referring now also to Fig. 12, a database may include medical data of diagnosed cases which can be linked to a medical diagnosis and/or categorized in other ways known in the art.

10 The medical data span a multi-dimensional space, including, for example, such parameters, or traits, as age, gender, patient/family history, test results (blood pressure, cholesterol, other essays), as well as vector quantities (EKG) and possibly also images. Fig. 12 is a two-dimensional diagram of two exemplary traits, trait 1(x-axis), trait 2 (y-axis) spanned by the database. Trait 1 may be a blood pressure reading and trait 2 an EKG characteristic signal. A point or area in the diagram will

15 have a likelihood of one or more medical diagnoses associated with it. Schematically indicated in Fig. 12 are the mean values  $D_1$ ,  $D_2$  and  $D_3$  for three different exemplary diagnoses as well as elliptical envelopes 1242, 1244, 1241 surrounding each mean value  $D_1$ ,  $D_2$  and  $D_3$  and representing an accepted confidence limit of value pairs (trait 1, trait 2) for a respective diagnosis  $D_1$ ,  $D_2$  and  $D_3$ . The mean positions  $D_1$ ,  $D_2$  and  $D_3$  of each cluster is unique to the particular diagnosis.

A first client may provide a server with exemplary patient data 1201 having a range for traits 1 and 2. The patient data 1201 overlap with both the envelope 1244 for diagnosis  $D_2$  and with envelope 1241 for diagnosis  $D_3$ . In other words, the client

25 could have either the disease  $D_2$  or the disease  $D_3$ . Likewise, a second client may provide the server with exemplary patient data having a range 1202 for the same traits 1 and 2. The patient data 1202 overlap with both the envelope 1242 for diagnosis  $D_1$  and with envelope 1244 for diagnosis  $D_2$ . In other words, the second client could have either the disease  $D_1$  or the disease  $D_2$ . Additional traits 3, ..., N

30 may help discriminate against the various possibilities for  $D_i$ .



As also seen from Fig. 12, the ranges 1201 and 1202 overlap in overlap region 1222. Moreover, a significant portion of overlap region 1222 also overlaps with the likelihood region for disease  $D_2$  as defined by envelope 1244. The system will hence group clients into a self-assembled group and associate them with the diagnosis  $D_2$ . In other words, the aforescribed method will automatically cluster the patients into finely granulated interest groups, in turn allowing users to easily obtain the richest information available concerning the details of their condition. Thus, the clustering of individuals is done objectively based on patients matching similar cases in the database, instead of by using subjective criteria.

A quantitative measure for the similarity or correspondence between data and/or data sets can be based on the Mahalanobis distance computed in a Mahalanobis metric. The Mahalanobis distance is a very useful way of determining the "similarity" of a set of values from an "unknown" sample to a set of values measured from a collection of "known" or reference data. Since the Mahalanobis distance is measured in units of standard deviations from the mean of the reference data, the reported matching values give a statistical measure of how well the unknown sample matches (or does not match) the reference data.

Visual inspection, like the diagram of Fig. 12, is usually not a viable method for actual discriminant analysis applications. An Euclidean distance method ("least-square-fit") does not take into account the variability of the values in the different dimensions, and is therefore not an optimum discriminant algorithm. The Mahalanobis distance, on the other hand, does take the sample variability into account. Instead of treating all values equally when calculating the distance from the mean point, it weights the differences by the range of variability in the direction of the sample point. Mahalanobis distances look at not only variations (variance) between the responses for the same trait, but also at the inter-trait variations (co-variance). The Mahalanobis group defines a multi-dimensional space whose boundaries determine the range of variation that are acceptable for unknown samples to be classified as members and hence admitted to the respective group.

Referring now to Fig. 13, in an exemplary process 1330 for forming/joining a self-assembled group and communicating anonymously with other members of the group, a user/patient connects to the server with the medical database, step 1331, and populates a datasheet with the patient's data, step 1332. As mentioned above, 5 populating the data sheet is not limited to entering text information. In many cases, the patient will respond to structured database queries (e.g., "What is your blood pressure?", "Have you experienced dizziness?", etc.); however, the patient may also wish to enter data or additional information in form of a comment, wherein the comment may not find a matching entry in the database, step 1333. Such comment 10 may be valuable and may lead to a future modification and/or addition to a "labeled" database entry in the event that more patients enter similar comments. One example may be hitherto unknown side effects in a clinical or drug study.

In the database, the data/comments entered by the patient are mapped onto the data of diagnosed or "labeled" cases, step 1334, and other users/patients mapped 15 onto the same data of diagnosed cases are identified, step 1335. At this point, the patient who entered the data/comments becomes a member of the same group and shares the diagnosis. The clients that are part of the group may have logged on anonymously or may have identified themselves to the server. In any event, the system is designed so that a client's identity is not revealed to another client, for 20 example, by assigning an ID to the clients. Having assigned a client ID, the server can now facilitate anonymous contact between clients that belong to the same group, step 1336.

Anonymous communication can be established between clients of a group via an anonymous link, for example, through email addresses that are never 25 associated with the user's identity or alternatively via a telephone patching system, where a third part establishes a one-time telephone connection without either of the two parties having any personal information about the other. When establishing anonymous email communication, email originating from one of the users can be sent to a third party with a unique, coded identifier corresponding to the 30 communication path between the two parties. The third party forwards the email

messages appropriately, and also allows either party to break the link at any time, in which the coded identifier becomes invalid, and all further communication between the two parties is terminated. Thus, the grouping communication system allows users to have open discussions with exchange of medical information and  
5 development of support networks, while still having the security that the identities of the users are protected and that users are in control of the email they receive.

The anonymous link can, in fact, be designed in such a way that the service providers themselves (i.e., the database, server and/or search engine), have no knowledge of true user identities. By storing encrypted markers associated with  
10 anonymous email accounts, but never connected to the identities of the individuals, the server assures that anonymity is maintained. When an anonymous communication channel is requested between two members of the same group, the system can, for example, establish the connection by passing the anonymous email addresses to a re-mailer.

15 As indicated in Fig. 14, the database entries and patient data entries 1432 do not only include text entries 1442 containing medical information, but can also accommodate full multimedia matching of any modality of available data. Similar to a physician combining the available information when assessing the similarity of two cases, the automatic matching system can use a combination of text 1442 (e.g.  
20 patient history), scalar measurements 1444 (e.g. blood pressure), vector quantities 1446 (e.g. EKG), and images 1448 (e.g. MR/CT).

Patient to patient support group development is one of the major types of communication channels facilitated by the automatic clustering techniques of the present invention. There are a variety of ways of structuring the actual initiation of  
25 communication between individuals. In one embodiment, patients whose information is in the database are never contacted by patients who enter their own information. When a patient enters his or her data and asks for a search, at the end of the process (or at some other convenient time, the search engine will ask the client two questions:

1. Would the client like to initiate communication with people who have similar matches to his (her) own information in a confidential manner, and

2. Would the client be willing to be contacted by others in a like confidential manner in order to discuss his (her) case with someone else.

5           Within a short period, a large number of users with varying diagnoses and assigned to different groups may have agreed to participate in an inter-person communication. The search engine and related software would keep track of the case numbers matching a specific group and would also keep track of the case numbers in the database for each client who agreed to communicate anonymously.

10       In this way, whenever a new client indicates that he/she would like to communicate with people being diagnosed with a similar disease, the software would search all clients who agreed to communicate with others to see if any of these other people's medical information matched the same cases in our database as the new client. The communication would be useful to the clients because each could share experiences

15       with various decisions, outcomes, treatments, etc. The initial communication with the server would be anonymous via email, keeping names and identities separate until both parties mutually agree to identify themselves.

          The described system and method may also allow patients to communicate with others having similar conditions, but without an associated diagnosis. In this

20       case, like in the previously discussed case, an anonymous communication link 16 can be established between a patient submitting a query and a patient whose medical information exists in the reference database. The process used throughout should not offend a participating institution that supplied the database, should not violate confidentiality/privacy issues, and should not harm in any way the patient or family

25       of the patient whose image was matched to the client, since the patient may not be aware of that his/her image was a part of the medical information search engine. As in the example above, a patient would indicate after using the search engine that he/she would like to communicate with a matching person from the database or the person's family. The trusted proprietors of the database would then send the request

30       to the institution from where the image was originally received, using the encoded

patient ID number, matching it to the list that is located in the computer bank of the sending institution. The sending institution would then search for the original patient's name, address, and e-mail if available, and would then either send a personalized form letter or e-mail to the patient or family. This letter can originate  
5 from the treating institution and explain that the treating institution is involved in this research. The patient or family will be told that the patient's medical image was included anonymously in a database of information. The letter or e-mail will explain briefly the purpose of the search engine and that an individual who used the system would like to contact the patient or family anonymously through the  
10 company with the database and search engine in order to have discussions about the patient's treatment, outcomes, etc. The patient or family member receiving the communication from the treating institution would be told that they were under no obligation to identify him(her)self to the inquiring client, and in fact did not have to reply to the letter if they did not wish to go any further. If the patient or family  
15 member agree to communicate, then they could reply via e-mail back to the sending institution, who would forward the e-mail to the proprietors of the database who would forward the communication back to the inquiring client. This would all be performed anonymously to protect the privacy of both the client and the patient or family.

20 The scope of the proposed automatic clustering system need not be restricted to groups of those patients who submitted their own medical information. The users of the system may also include other individuals and institutions associated with the health care industry, such as radiologists, physicians, clinical researchers, and pharmaceutical companies. In addition to the patient-patient interactions facilitated  
25 by the clustering system, other combinations of interactions between individuals are made possible, such as patient-physician, patient-researcher, physician-radiologist, etc. For example, a clinical researcher investigating multiple sclerosis (MS) is likely to submit many medical images of patients afflicted with the condition. When an MS patient submits his or her own scan, and it matches other scans containing MS

lesions, the patient and clinical researcher will be grouped together, opening a new communication channel.

Referring now to Fig. 15, the disclosed system and process 1550 can also be used in a clinical trial setting, where patients searching for alternative treatment options for their disease or condition may consider turning to clinical trials of experimental drugs or treatments. The medical image database coupled with the self-organizing groups provides a means of connecting these patients with institutions performing clinical trials. The patient or hospital connect to the server, step 1551, and submits a patient's medical information, e.g., a scan and additional data, to the database, step 1552, where the matching medical images are returned with possible diagnoses 1553. Based on the matching diagnoses, it is checked if an appropriate clinical trial is currently registered with the database, step 1554. If the answer is affirmative, the patient or physician has the option of initiating a communication path with the investigators of the clinical trials. Upon consent, the clinical researchers would have anonymous access to the medical images of the patient interested in clinical trials, step 1555, in order to evaluate whether the patient matches the requirements for the trial. By matching the patient's medical images to the database, clinical researchers interested in a certain condition and patients afflicted with that condition seeking care alternatives are automatically clustered into the same group, step 1556. This approach can significantly broaden the number of candidates that can be enlisted in a clinical trial.

The medical information submitted by the users who were grouped together include uninterpreted as well as interpreted information. Interpreted information refers to information in the database that has been "labeled" or associated with a diagnosis. Conversely, uninterpreted information include patient information that has not been associated with a diagnosis, but may also include information supplied by the patient as a comment and for which no entry exists yet in the database. Referring back to Fig. 12, uninterpreted information is represented by the areas of 1201 and 1202 that do not overlap with any of the exemplary diagnoses  $D_1$ ,  $D_2$  and  $D_3$ . However, such uninterpreted information can be useful for forming groups and

associating users with existing groups. For example, the point (X) in Fig. 12 is shared by multiple users and can be used to form a group that includes two users, even if a diagnosis has not yet been established. Such uninterpreted information that can originate from user-supplied comments, can be stored in a "Watch Database" at the server, as indicated as step 1666 in process 1660 depicted in Fig. 16.

As described above with reference to Fig. 13, a user/patient may input in step 1633 additional comments, which can be text, measurement data and/or image information, that could be provided as an answer to a question: "Do you have any other symptoms?" Referring now back to Fig. 16, the system will check if language or features in the comment can be associated with "labeled" information in the database, step 1662. If this is the case, then the user's comments are matched up with database entries, step 1634. If the database has no corresponding entries, then the databases, including the "Watch Database", are searched for similar comments from other users, step 1664. If no similar comments are found, the user's comment is entered into the "Watch Database" for future comparisons, step 1666. If corresponding entries are found in step 1664, then it appears likely that the hitherto unlabeled feature is important in context of a medical diagnosis, and the process 1660 may request or at least suggest that this unlabeled feature be included and "labeled" in the database, step 1666. Techniques for arriving at criteria for two features (text, images, etc.) to be viewed as being similar are known in the art.

In summary, grouping invokes a bootstrapping technique where the initial collection of information has been labeled or interpreted, whereas the submitted queries are usually "unlabeled". Once two unlabeled queries are matched to similar labeled information in the database, these unlabeled queries are interpreted, and a group is formed based on both the labeled and the originally unlabeled information. The automatic grouping approach provides a multitude of advantages to existing subjective methods of forming interest groups that rely heavily on the individuals being able to precisely and accurately express exactly where their interests lie.

Thus, while the invention has been disclosed in connection with the preferred embodiments shown and described in detail, various modifications and

improvements thereon will become readily apparent to those skilled in the art. It should be understood that all matter contained in the above description or shown in the accompanying drawings shall be interpreted as illustrative, and not in a limiting sense, and that the following claims should be interpreted in the broadest sense

5 allowable by law.



**Claims:**

1. A method comprising:
  - receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;
  - 5 registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of registered images;
  - extracting one or more feature vectors from each one of the plurality of
  - 10 registered images, each one of the one or more feature vector including data derived from image data of one of the plurality of registered images;
  - receiving a query image, the query image including an instance of a human body part obtained through the medical imaging technique;
  - registering the query image to the coordinate system in a non-rigid manner to
  - 15 obtain a registered query image;
  - identifying one or more regions of interest in the registered query image;
  - extracting one or more feature vectors from the one or more regions of interest in the registered query image; and
  - matching the query image to one or more of the plurality of images by
  - 20 comparing the one or more feature vectors from the one or more regions of interest in the registered query image to one or more feature vectors of one or more of the plurality of registered images.
2. A method comprising:
  - 25 receiving a plurality of images obtained using magnetic resonance imaging, each one of the plurality of images including a planar slice of a human head;
  - registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more features of each one of the plurality of images that is a planar slice from a specified z-position with a corresponding one
  - 30 or more features of each other one of the plurality of images that is a planar slice

from the specified z-position, thereby obtaining a plurality of registered images, registering including at least one of normalizing, rotating, or warping each one of the plurality of images;

5        sampling each one of the plurality of registered images in a plurality of non-overlapping tiles;

             extracting one or more feature vectors from each one of the non-overlapping tiles for each one of the plurality of registered images, each one of the one or more feature vector including at least one of a mean signal intensity, a standard deviation of signal intensity, a mean edge magnitude, a mean x-direction gradient, a mean y-  
10        direction gradient, a mean absolute value of surface curvature, and a mean absolute value of levelset curvature;

             receiving a query image, the query image obtained using magnetic resonance imaging and including a planar slice of a human head of a patient;

             registering the query image, sampling the query image, and extracting one or  
15        more feature vectors from the query image as performed above for the plurality of images;

             identifying one or more regions of interest in the registered query image by selecting a point within the query image; and

             matching the query image to one or more of the plurality of images by  
20        comparing the one or more feature vectors from the one or more regions of interest in the registered query image to one or more feature vectors of one or more of the plurality of registered images.

3.        The method of claim 1 or 2 wherein the one or more regions of interest in the  
25        registered query image are located inside and/or outside the one or more regions of interest, or can include the entire registered query image.

4.        The method of claims 1-3 wherein matching the query image to one or more  
of the plurality of images is automatically scheduled based upon data in a patient  
30        record.

5. The method of claims 1-4 further comprising providing a score indicative of a similarity between the query image and each one of the plurality of images.
- 5 6. The method of claims 1-5 further comprising identifying one or more regions of interest in each one of the plurality of images.
7. The method of claim 6 further comprising matching the one or more feature vectors from the one or more regions of interest in the registered query image to the  
10 one or more regions of interest in each one of the plurality of registered images.
8. The method of claims 1-7 further comprising identifying a relative z-position in each one of the plurality of registered images, the relative z-position corresponding to a position of a planar image along a z-axis.
- 15 9. The method of claim 8 further comprising training a model to determine a relative z-position of the registered query image.
10. The method of claims 1-9 further comprising filtering the plurality of  
20 registered images according to one or more filtering criteria to obtain a subset of the plurality of registered images to match to the registered query image.
11. The method of claims 1-10 further comprising associating a pathology with the query image based upon pathology data for each of the plurality of registered  
25 images matched to the registered query image.
12. The method of claims 1-11 wherein the plurality of images include at least one of magnetic resonance images or computerized tomography images representing at least one of head images, neck images, spine images, or musculo-skeletal images..

30

13. The method of claims 1-12 further comprising associating a pathology with the query image based upon the matched feature vectors and one or more known pathologies associated with the plurality of images.

5 14. The method of claims 1-13 wherein registration comprises at least one of normalizing, masking, scaling, rotating, or warping and/or dividing each one of the plurality of registered images into a plurality of discrete tiles and sampling each one of the plurality of registered images within each one of the discrete tiles.

10 15. A computer program product comprising:

computer executable code for receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;

15 computer executable code for registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of registered images;

20 computer executable code for extracting one or more feature vectors from each one of the plurality of registered images, each one of the one or more feature vector including data derived from image data of one of the plurality of registered images;

computer executable code for receiving a query image, the query image including an instance of a human body part obtained through the medical imaging technique;

25 computer executable code for registering the query image to the coordinate system in a non-rigid manner to obtain a registered query image;

computer executable code for identifying one or more regions of interest in the registered query image;

30 computer executable code for extracting one or more feature vectors from the one or more regions of interest in the registered query image; and

computer executable code for matching the query image to one or more of the plurality of images by comparing the one or more feature vectors from the one or more regions of interest in the registered query image to one or more feature vectors of one or more of the plurality of registered images.

5

16. The computer program product of claim 15 further comprising computer executable code for identifying one or more regions of interest in each one of the plurality of images.

10 17. The computer program product of claim 16 further comprising computer executable code for matching the one or more feature vectors from the one or more regions of interest in the registered query image to the one or more regions of interest in each one of the plurality of registered images.

15 18. The computer program product of claim 15 further comprising computer executable code for identifying a relative z-position in each one of the plurality of registered images, the relative z-position corresponding to a position of a planar image along a z-axis.

20 19. The computer program product of claim 18 further comprising computer executable code for training a model to determine a relative z-position of the registered query image.

20. The computer program product of claim 15 further comprising computer  
25 executable code for filtering the plurality of registered images according to one or more filtering criteria to obtain a subset of the plurality of registered images to match to the registered query image.

21. A system comprising:

an image database connected to a network that stores a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique, and each one of the plurality of images pre-processed by registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of registered images, each one of the plurality of images further pre-processed to extract one or more feature vectors from each one of the plurality of registered images, each one of the one or more feature vector including data derived from image data of one of the plurality of registered images;

a client device connected to the network and configured to receive a query image including an instance of a human body part obtained through the medical imaging technique, the client device registering the query image to the coordinate system in a non-rigid manner to obtain a registered query image, and the client device receiving a user input that identifies one or more regions of interest in the registered query image, the client device further configured to extracting one or more feature vectors from the one or more regions of interest in the registered query image; and

a server connected in a communicating relationship with the image database and the client device, the server configured to match the query image to one or more of the plurality of images by comparing the one or more feature vectors from the one or more regions of interest in the registered query image to one or more feature vectors of one or more of the plurality of registered images, the server further configured to provide one or more matched images to the client device.

25

22. The system of claim 21, the server further configured to provide to the client device clinical data associated with each of the one or more matched images.

23. A method comprising:

receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;

registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of registered images;

labeling each one of the plurality of registered images with a label according to an observed characteristic;

obtaining one or more feature vectors from each one of the plurality of registered images; and  
training a model to associate the one or more feature vectors with the label.

24. The method of claim 23 further comprising applying the model to generate a label for each one of an additional plurality of images.

25. The method of claim 23 wherein each label is a position of an image on a z-axis perpendicular to an image plane of the image.

26. The method of claim 23 wherein each label is a region of interest.

27. The method of claim 23 wherein each label is at least one of an age, a sex, a diagnosis, a presence of contrast agents, an image type, or a diagnostic significance of a region of interest.

28. The method of claim 23 wherein the model is derived from a statistical learning methodology.

29. The method of claim 23 wherein the model is a linear regression model.

30. The method of claim 23 wherein partial least squares are used to determine one or more coefficients of the model.

31. The method of claim 23 further comprising applying the model to generate  
5 labels for a second database of images.

32. The method of claim 23 further comprising:

locating an image database that is accessible through a network, the image database including a second plurality of images;

10 registering each one of the second plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the second plurality of images within the coordinate system, thereby obtaining a second plurality of registered images;

applying the model to label each one of the second plurality of images; and

15 searching the image database using the labels to evaluate a similarity of a query to one or more of the second plurality of images.

33. The method of claim 23 further comprising organizing a plurality of databases by locating images, registering images, and labeling images within each  
20 one of the plurality of databases, each one of the plurality of databases being labeled with a different model, and searching the plurality of databases using the labels to evaluate a similarity of a query to one or more records in each of the plurality of databases.

25 34. A computer program product comprising:

computer executable code for receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;

30 computer executable code for registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like



features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of registered images;

computer executable code for labeling each one of the plurality of registered images with a label according to an observed characteristic;

5 computer executable code for obtaining one or more feature vectors from each one of the plurality of registered images; and

computer executable code for training a model to associate the one or more feature vectors with the label.

10 35. The computer program product of claim 34 further comprising computer executable code for applying the model to generate a label for each one of an additional plurality of images.

36. A method comprising:

15 receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;

registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of

20 registered images;

receiving a header for each one of the plurality of images that includes data associated with the one of the plurality of images;

obtaining one or more feature vectors from each one of the plurality of registered images;

25 training a model to associate the one or more feature vectors with the header; and

applying the model to identify the presence of any errors in a new header for a new image.

37. The method of claim 36 wherein the header includes a magnetic resonance imaging characteristic.

38. The method of claim 36 wherein the header includes at least one of a contrast agent attribute that indicates a presence or absence of a contrast agent, or a sequence type attribute that indicates a sequence type for the plurality of images, the sequence type being at least one of MRA or T1.

39. A method comprising:  
10 receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;  
registering each one of the plurality of images in a non-rigid manner to a coordinate system to superimpose one or more like features within each one of the plurality of images within the coordinate system, thereby obtaining a plurality of  
15 registered images;  
obtaining one or more feature vectors from each one of the plurality of registered images;  
associating a pathology with each one of the plurality of images;  
training a model to associate the pathology associated with each image with  
20 the one or more feature vectors for that image; and  
applying the model to identify the presence of the pathology in a new image.

40. The method of claim 39 further comprising obtaining each one of the one or more feature vectors from a region of interest within one of the plurality of images.

41. A method comprising:  
receiving a plurality of images, each one of the plurality of images including an instance of a human body part obtained through a medical imaging technique;  
registering each one of the plurality of images in a non-rigid manner to a  
30 coordinate system to superimpose one or more like features within each one of the

plurality of images within the coordinate system, thereby obtaining a plurality of registered images;

identifying one or more regions of interest in each of the plurality of registered images, the regions of interest including a pathology; and

5 generating a spatial probability map of locations of the pathology from the plurality of registered images and the regions of interest.

42. The method of claim 41 further comprising using the spatial probability map as a medical diagnostic aid.

10

43. A search engine comprising:

a first search mode in which a database of images is searched for one or more images similar to a query image;

a second search mode in which the database of images is searched for one or  
15 more images matching a location in an atlas, the atlas providing a three-dimensional depiction of an object; and

a third search mode in which the database of images is searched for a  
diagnosis that has associated therewith images similar to the query image, the  
diagnosis and a similarity score being provided for each image retrieved from the  
20 database of images.

44. The search engine of claim 43 wherein the first search mode and the third  
search mode employ a matching system in which a database of medical image  
studies is pre-processed to non-rigidly register one or more images of the medical  
25 image studies in the database to a global coordinate system to provide a plurality of  
registered medical images, and to extract features from the registered medical  
images, the query image being non-rigidly registered to the global coordinate system  
to provide a registered query image, and one or more feature vectors being extracted  
from the registered query image, a match to the database being performed by

comparing the feature vectors extracted from the plurality of registered medical images to the feature vectors extracted from the registered query image.

45. The search engine of claim 44 wherein one or more regions of interest are provided for each one of the registered medical images and one or more regions of interest are provided for the registered query image, the match to the database being performed only for those registered medical images having at least one region of interest coinciding with at least one of the one or more regions of interest for the registered query image within the global coordinate system.

10

46. A method for radiology workflow management comprising:  
displaying in a first window of a computer display a magnetic resonance image study of a head to a physician in a user-navigable electronic form;  
displaying in a second window of a computer display one or more magnetic resonance images matched to the magnetic resonance image study of the head by a computerized matching system;

15

displaying in a third window of a computer display a multi-media reference for abnormalities associated with heads, the multi-media reference including text and images;

20

displaying in a fourth window of a computer display one or more images matched to at least one of the magnetic resonance image study of the head or one of the text or images of the multi-media reference;

25

providing a user interface for navigating between the first window, the second window, the third window, and the fourth window, at least one of the first window, the second window, the third window, and the fourth window being displayed at any one time, the user interface further adapted to receive a clinical evaluation of the magnetic resonance image study.

47. The method of claim 46 further comprising displaying clinical data associated with the magnetic resonance image study.

30

48. The method of claim 46 further comprising displaying clinical data associated with the one or more images displayed in the fourth window and matched to at least one of the magnetic resonance image study of the head or one of the text or images of the multi-media reference.

49. The method of claim 46 wherein the first window includes an input for receiving an identification of one or more regions of interest within the magnetic resonance image study of the head and an input for initiating a search for matches to the magnetic resonance image study of the head, the search using the identified one or more regions of interest as a search criterion.

50. The method of claim 46 wherein the first window displays more than one magnetic resonance image study to a user for evaluation.

51. The method of claim 46 wherein the fourth window displays images matched from a remote database.

52. The method of claim 46 wherein the first window, the second window, the third window, and the fourth window are provided on a client device connected in a communicating relationship with a network.

53. The method of claim 52 wherein the first window, the second window, the third window, and the fourth window are at least one of a desktop application running on the client device, an application running on a remote service and provided to the client device as an application service, or an application running from a Web browser at the client device.

54. A method for radiology workflow management comprising:  
displaying in a first window of a computer display a magnetic resonance image study of a head to a physician in a user-navigable electronic form;

displaying in a second window of a computer display one or more magnetic resonance images matched to the magnetic resonance image study of the head by a computerized matching system;

5 displaying in a third window of a computer display a multi-media reference for abnormalities associated with heads, the multi-media reference including text and images;

displaying in a fourth window of a computer display one or more images matched to at least one of the magnetic resonance image study of the head or one of the text or images of the multi-media reference;

10 providing a user interface for navigating between the first window, the second window, the third window, and the fourth window, at least one of the first window, the second window, the third window, and the fourth window being displayed at any one time, the user interface further adapted to receive a clinical evaluation of the magnetic resonance image study.

15

55. A method of assembling a group of at least two entities having a common interest, comprising

- (a) comparing data provided by a first entity to stored reference data,
- 20 (b) establishing a first correspondence between the data provided by the first entity and the reference data;
- (c) comparing data provided by at least one other entity to the stored reference data,
- (d) establishing at least one other correspondence between the data provided by  
25 the at least one other entity and the reference data;
- (e) evaluating a similarity measure between the first correspondence and the at least one other correspondence, and
- (f) forming the group of the first entity and the at least one other entity if the similarity measure exceeds a predetermined threshold.

30

1/15




Fig. 1

2/15




Fig. 2






Fig. 3

4/15

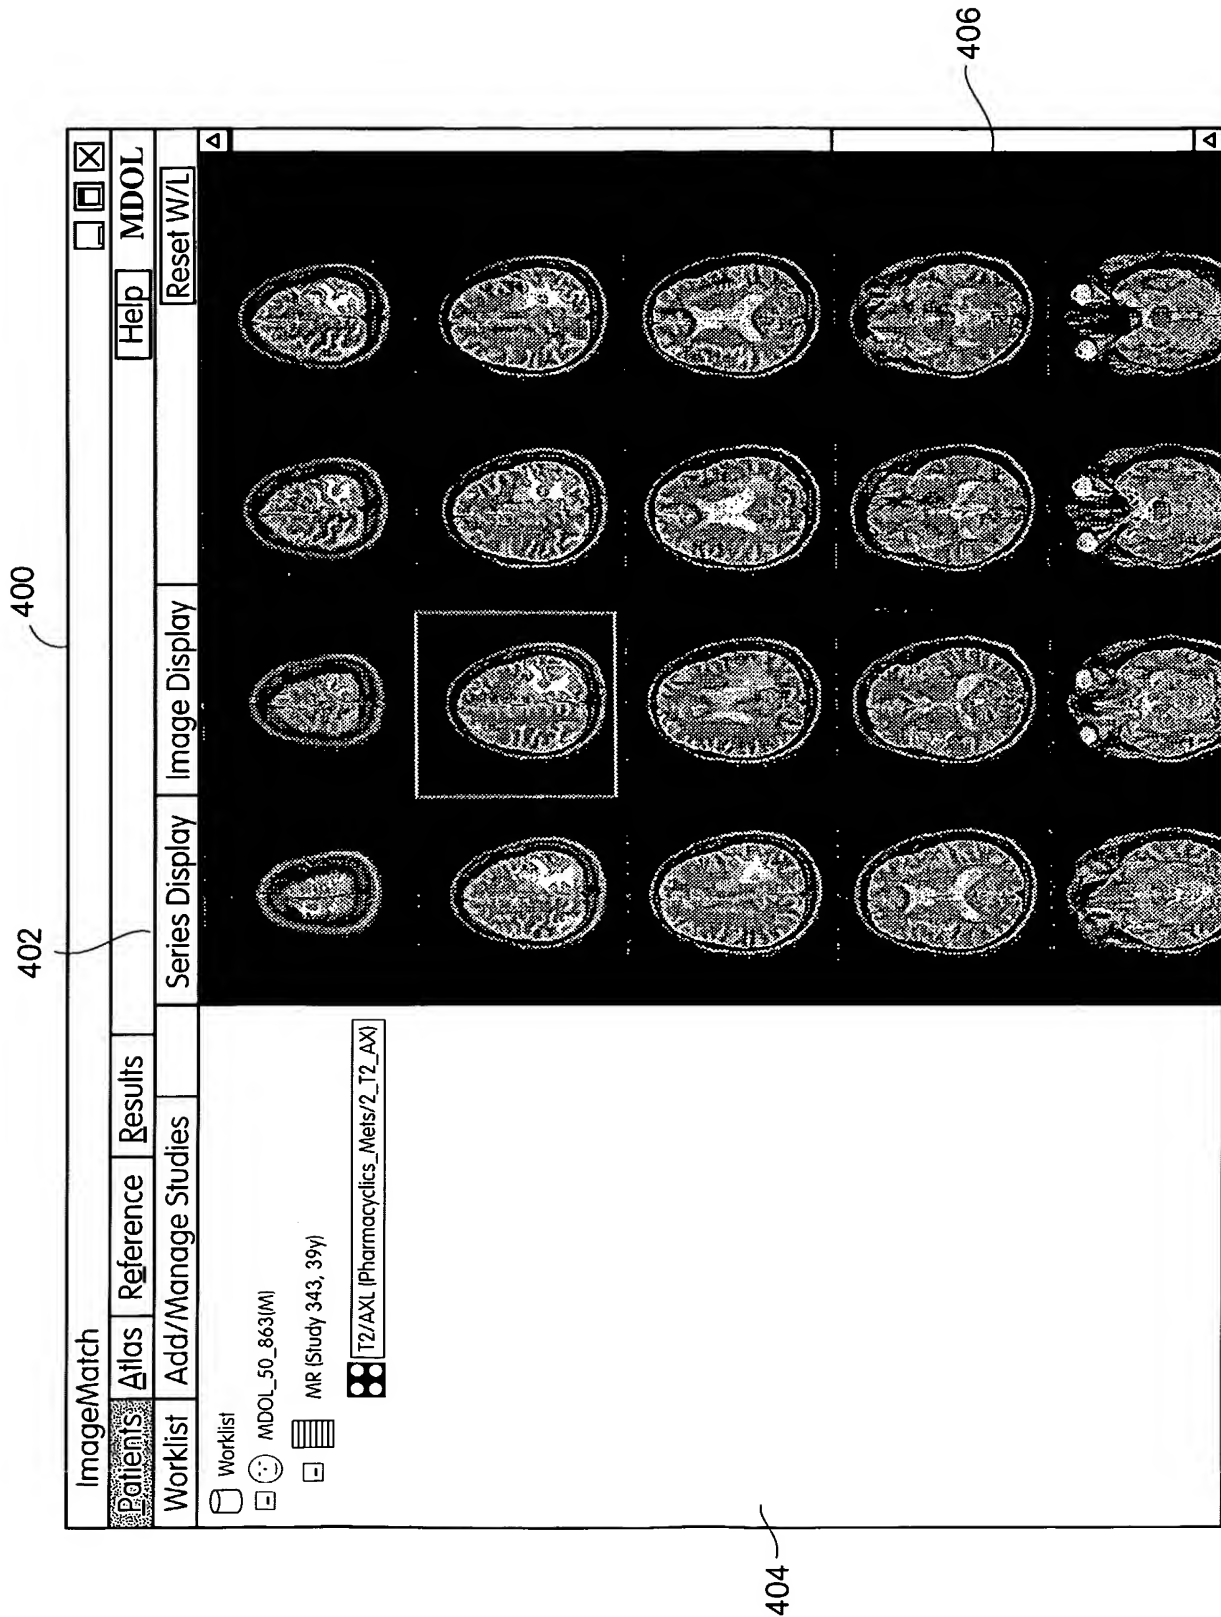
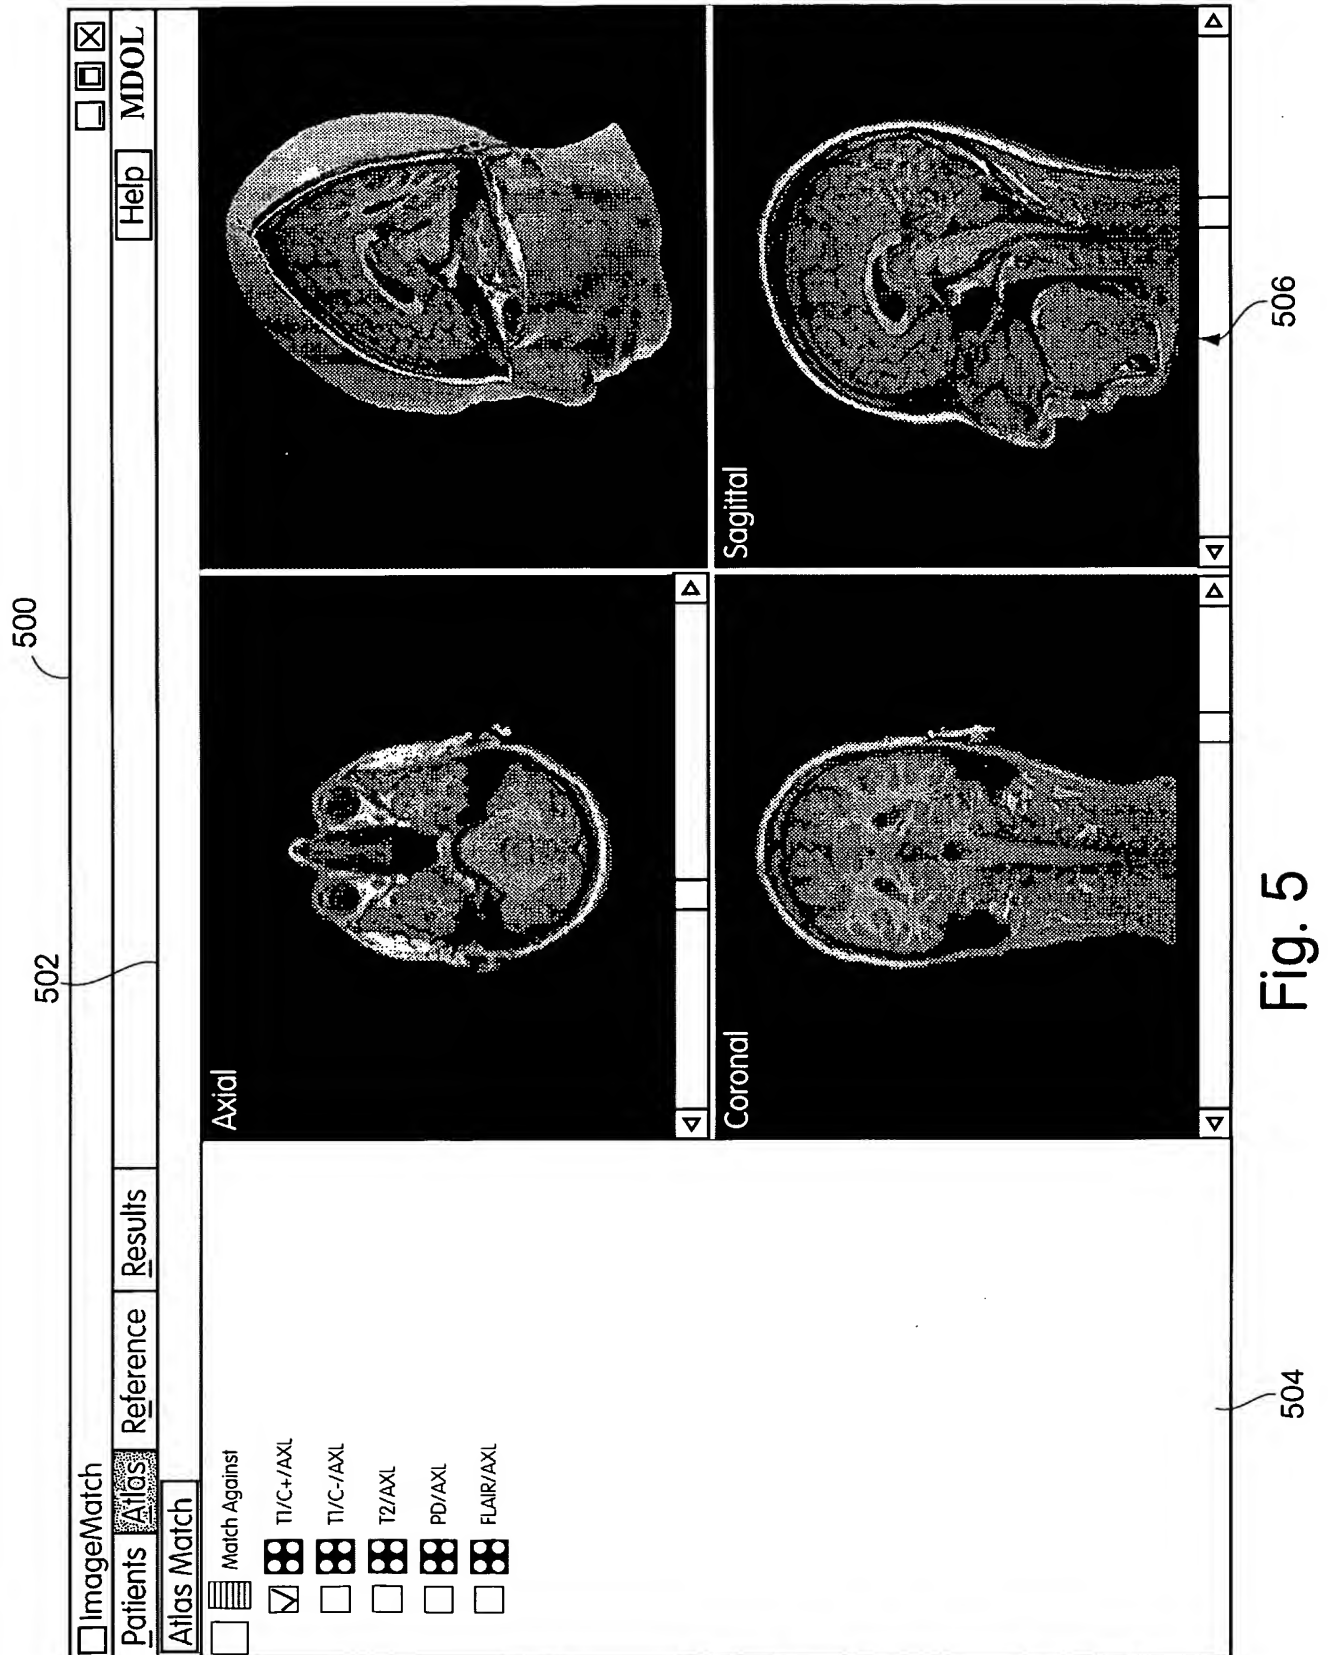
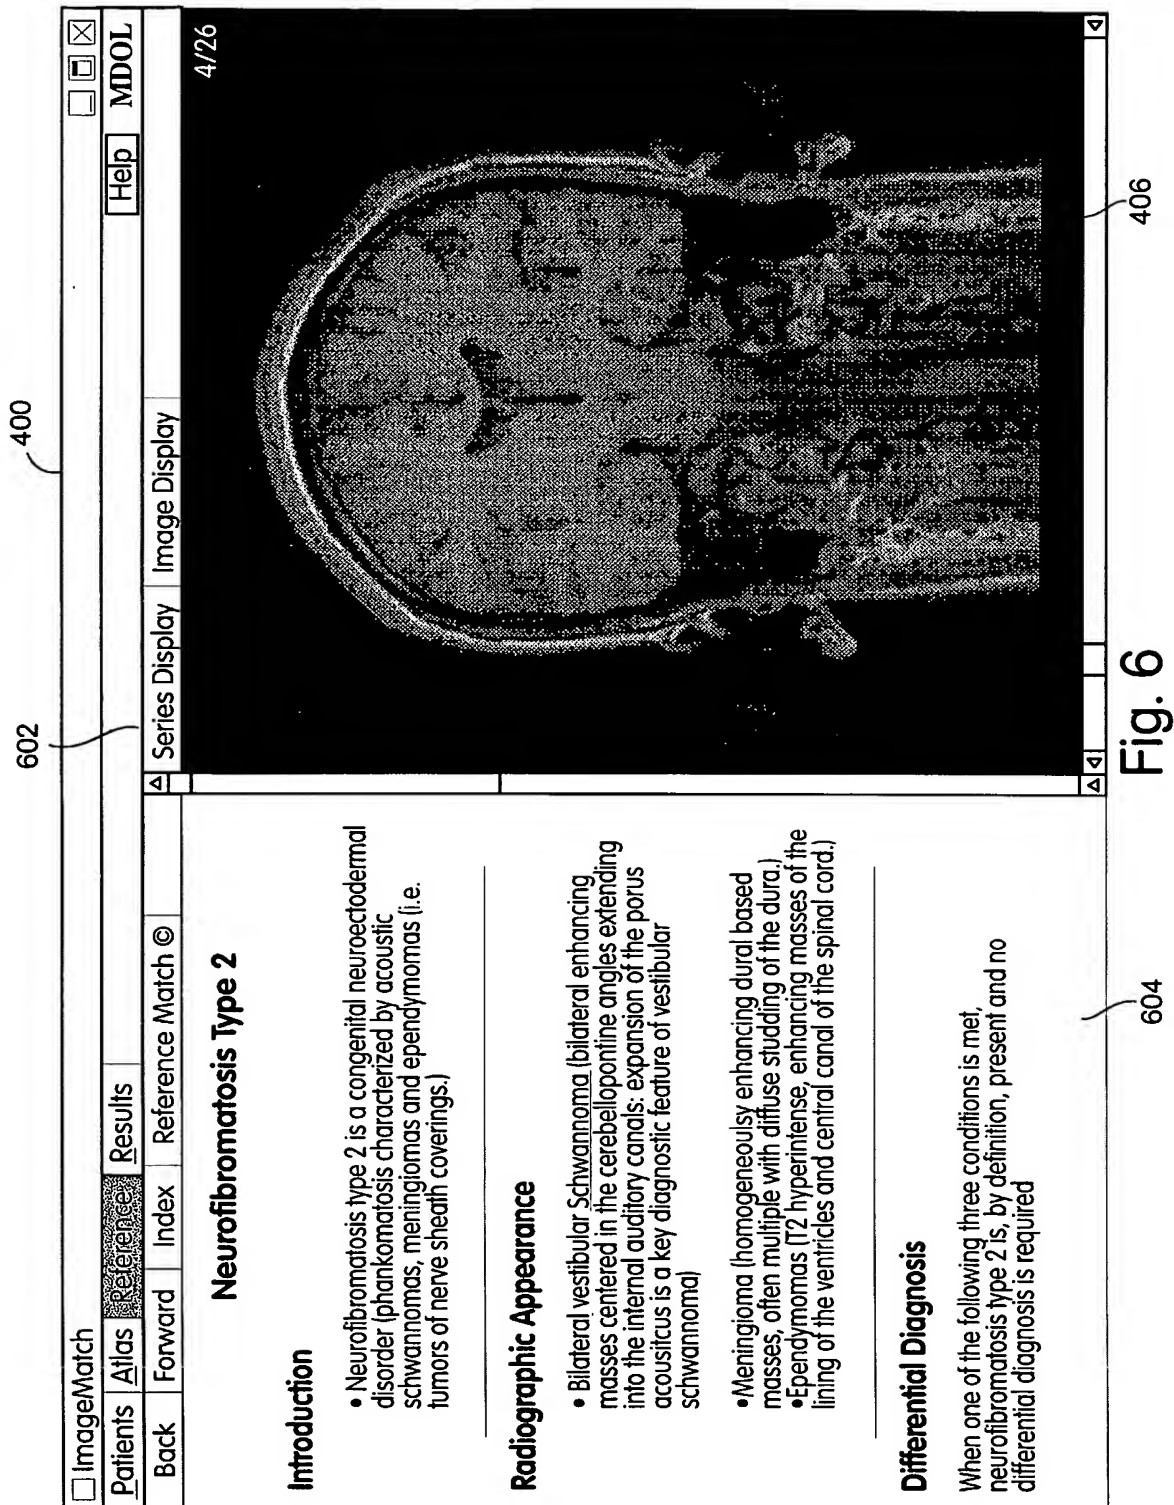


Fig. 4





7/15

ImageMatch
MDOL

Patients
Atlas
Reference
Results

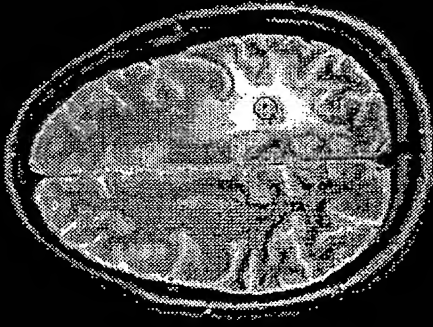
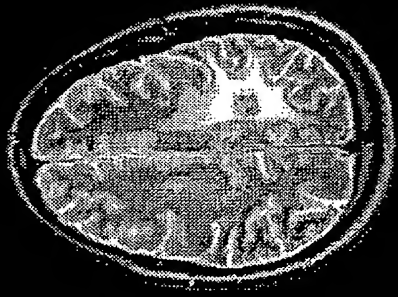
MDOL\_50\_863 (M)

MR (Study 343, 39y)

T2/AXL (Pharmacoclycs\_Mets/2\_T2\_AXI)

Query Series
Query Image
Results Series
Results Image

39/M
8/28
39/M
8/28

Results by Pathology
Available Series

☐ All Match Results (100 matches)

☒ INFECTION/INFLAMMATORY (55 matches)

☒ VASCULAR DISEASE (31 matches)

☒ INTRACRANIAL NEOPLASMS (10 matches)

☒ OTHER (4 matches)

Matching Images
Available Series
Reference
Clinical Report



METASTASES  
39/M/Not Verified  
(0.00852188, rank 1)



METASTASES  
38/M/Not Verified  
(1.37028, rank 2)



MULTIPLE SCLEROSIS  
53/F/Not Verified  
(2.03312, rank 3)



MULTIPLE SCLEROSIS  
48/F/Not Verified  
(2.10408, rank 4)

Fig. 7




Fig. 8

9/15




Fig. 9

10/15

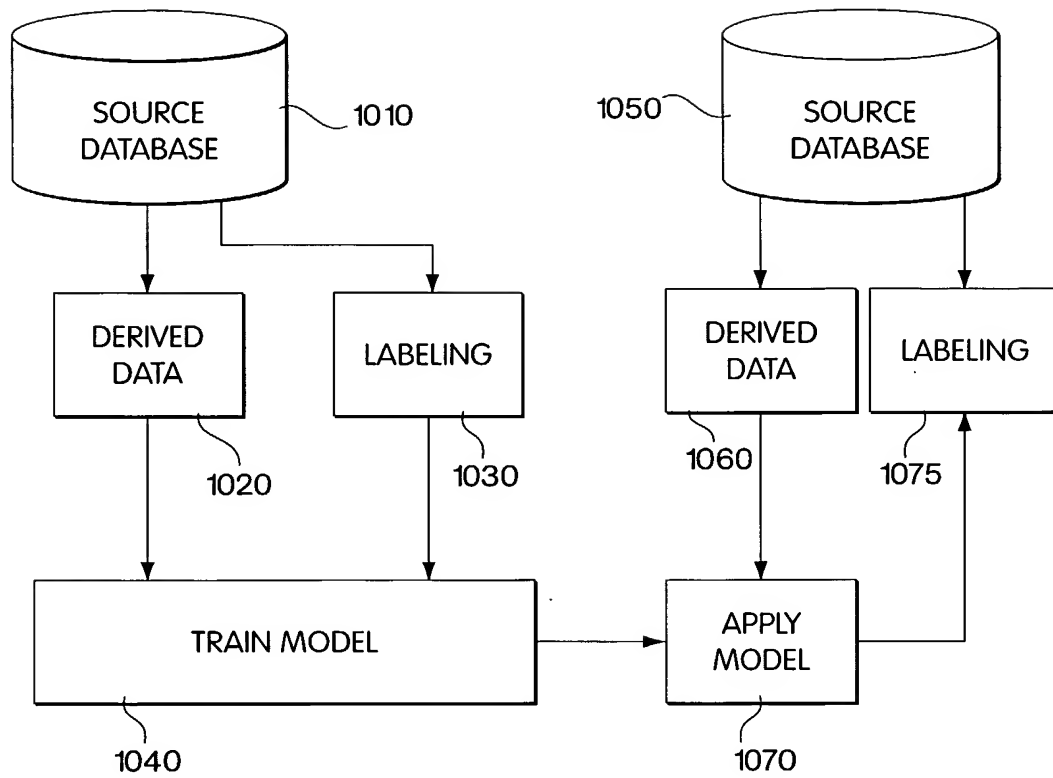


Fig. 10



11/15




Fig. 11

12/15




Fig. 12

13/15

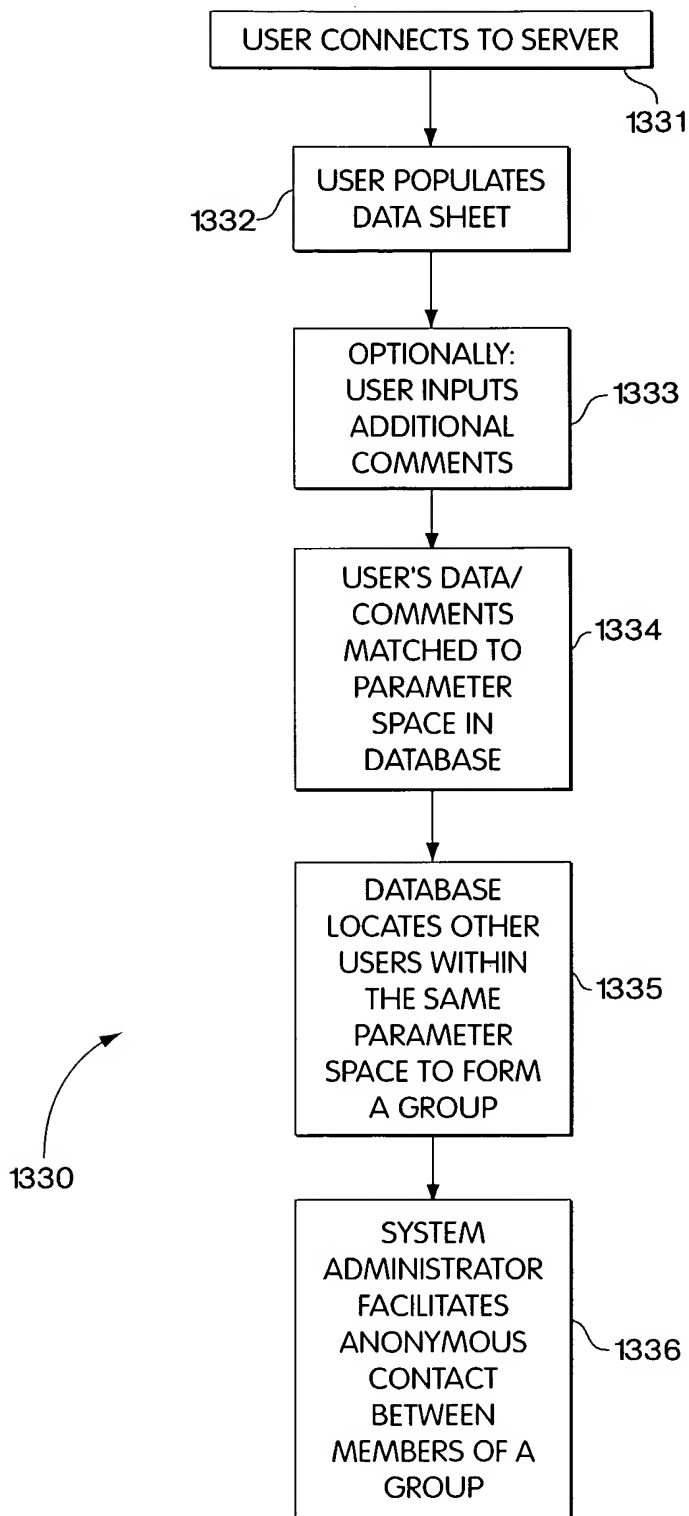


Fig. 13

14/15

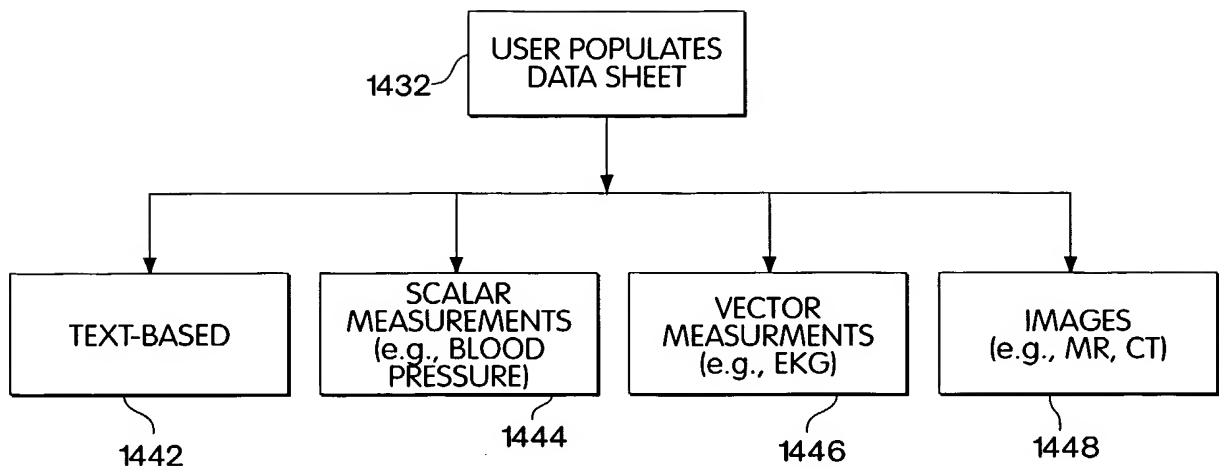


Fig. 14

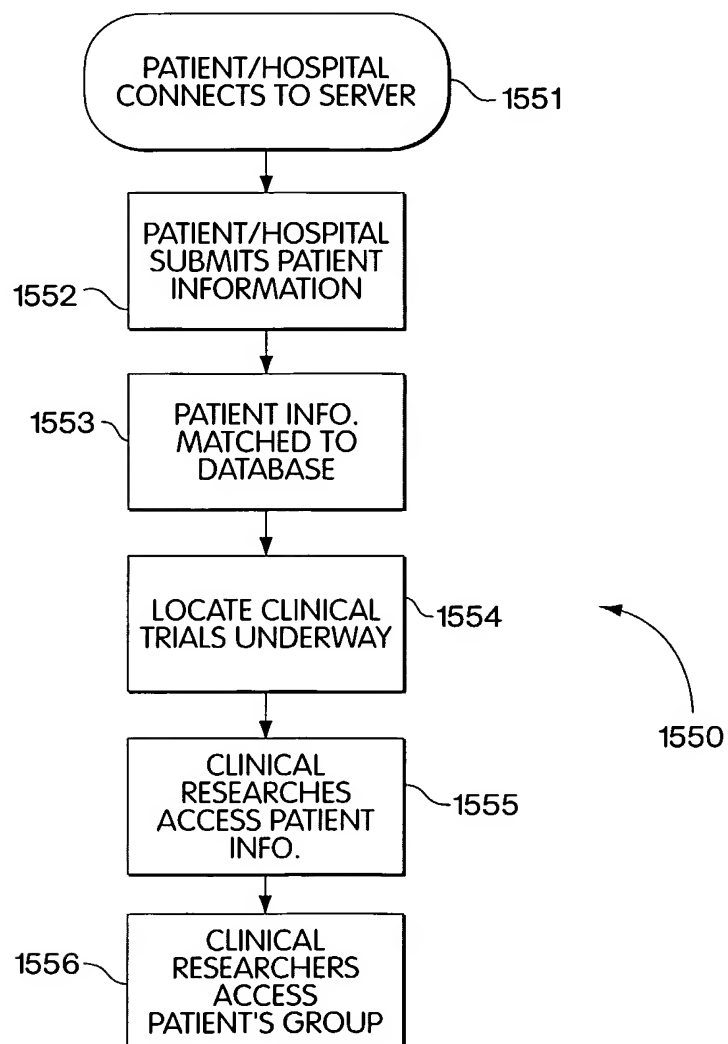


Fig. 15

15/15

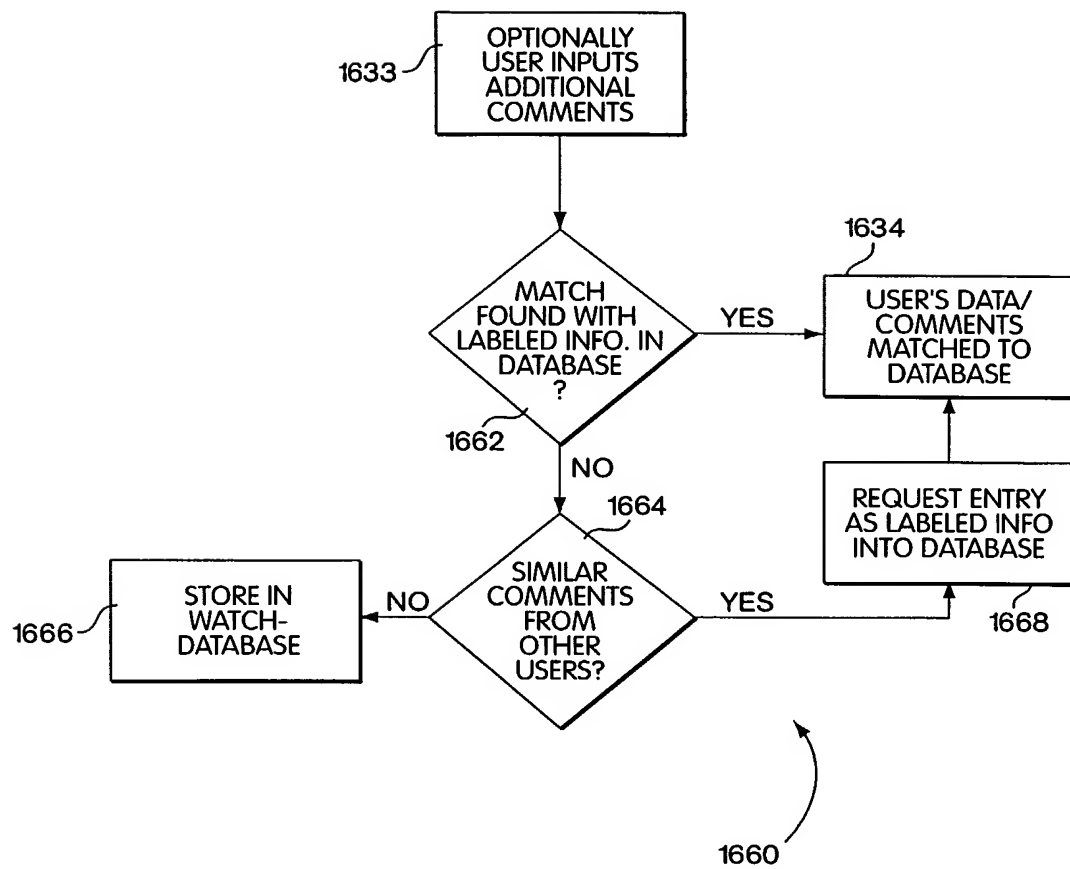


Fig. 16

## INTERNATIONAL SEARCH REPORT

Intern I Application No  
PCT/US 01/29573

A. CLASSIFICATION OF SUBJECT MATTER  
IPC 7 G06F17/30 G06T7/00

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
IPC 7 G06F G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

WPI Data, PAJ, INSPEC, EPO-Internal, BIOSIS

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	THURFJELL L ET AL: "A volumetric model for identifying anatomical structures in tomographic brain images" PROCEEDINGS OF THE 8TH SCANDINAVIAN CONFERENCE ON IMAGE ANALYSIS, PROCEEDINGS OF SCANDINAVIAN CONFERENCE ON IMAGE ANALYSIS, TROMSO, NORWAY, 25-28 MAY 1993, pages 473-478 vol.1, XP008009978 1993, Tromso, Norway, NOBIM-Norwegian Soc. Image Process. & Pattern Recognition, Norway abstract page 473, right-hand column, line 24 - line 28 page 476, right-hand column, line 37 - line 40 page 473, right-hand column, line 31 - line 32	1,8,15, 16,18
Y	-/--	5

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

## \* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

19 November 2002

Date of mailing of the international search report

06. 02. 2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

CHATEAU J.P.

## INTERNATIONAL SEARCH REPORT

Intern Application No  
PCT/US 01/29573

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A		2,6,7,9, 17,19,20
A	--- WO 00 17751 A (MACK WALTER J ;INTEL CORP (US); KORTA JEAN CHARLES (US)) 30 March 2000 (2000-03-30) abstract	1
A	--- FERRANT M ET AL: "MULTI-OBJECT SEGMENTATION OF BRAIN STRUCTURES IN 3D MRI USING A COMPUTERIZED ATLAS" PROCEEDINGS OF THE SPIE, SPIE, BELLINGHAM, VA, US, vol. 3661, February 1999 (1999-02), pages 986-995, XP000892602 page 994, line 22 - line 24	11-13
Y	--- EP 0 947 937 A (CANON KK) 6 October 1999 (1999-10-06) abstract -----	5

## INTERNATIONAL SEARCH REPORT

Int'l application No.  
PCT/US 01/29573

### Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:  
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

### Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-20

#### Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☐ No protest accompanied the payment of additional search fees.



FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. Claims: 1-20

A method comprising receiving a plurality of images, registering each of the plurality of images to a coordinate system, receiving a query image, registering this query image to the coordinate system and matching the query image to one or more of the plurality of images by comparison of one or more feature vectors of each image

2. Claims: 21-22

A client- server system

3. Claims: 23-35

A method of training a model to associate labels with image features

4. Claims: 36-38

A method for training a model for the association of a header to image features and applying the model to identify the presence of any errors in a new header

5. Claims: 39-40

A method of training a model to associate a pathology to one or more features of an image and applying the model to identify the presence of a pathology in a new image

6. Claims: 41-42

A method for generating a spatial probability map of location of a pathology in regions of interest identified in a set of images.

7. Claims: 43-45

A search engine

8. Claims: 46-54

A method of radiology workflow management

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

9. Claim : 55

A method of assembling a group of at least two entities  
having a common interest

## INTERNATIONAL SEARCH REPORT

Intern I Application No  
PCT/US 01/29573

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0017751	A	30-03-2000	US 6178392 B1	23-01-2001
			AU 6039099 A	10-04-2000
			DE 19983574 T0	16-08-2001
			GB 2357949 A ,B	04-07-2001
			JP 2002525746 T	13-08-2002
			TW 405059 B	11-09-2000
			WO 0017751 A1	30-03-2000
			US 6505148 B1	07-01-2003
-----				
EP 0947937	A	06-10-1999	JP 11288418 A	19-10-1999
			JP 11316819 A	16-11-1999
			EP 0947937 A2	06-10-1999
-----				